STATE OF NEW JERSEY

SHARED IT ARCHITECTURE

Executive Summary

The purpose of this document is to guide Executive Branch Agencies toward leveraging existing shared IT infrastructure, processes and support staff in order to minimize risk and lower the overall cost of IT projects.

This document focuses on the existing shared infrastructure used by multiple State agencies and is not a complete listing of every product used by every State agency.

The State's Shared IT Infrastructure has been built to support this vision. It is a robust, standardized environment that currently supports Executive Branch computer systems within and across agency boundaries. The infrastructure is designed to rapidly accommodate growth and replacement of hardware, middleware, software and communications as new business needs arise or when efficiencies can be realized by upgrading or replacing existing components.

It is also the intent of Executive Order #42 to consolidate agency operations and eliminate redundant functions in order to provide the best quality of service, most efficient use of staff and computer space, reduce energy consumption, and achieve the flexibility required to maintain a state of the art technology environment to meet the needs for delivering services to the State's residents, employees and business partners.

This document is intended to provide sufficient technical detail regarding the various components of the State's Shared IT Infrastructure and, in Appendix 1, denotes the level of support and investment the State has made in specific products and technologies. While continually evolving, it is based on industry standard open system solutions that provide a high degree of vendor neutrality, maximum flexibility, and the agility needed to meet the ever-growing service delivery needs of the State's Executive Branch. The use of open standards is critical to the State's ability to interact with constituents and business partners across the internet. The focus on specific products and technologies is equally important in order to minimize the staffing resources needed to support a shared, consolidated infrastructure.

The organization of this document is based on the IT Architecture Stack depicted below, where each layer represents a set of technologies put in place to support specific business processes. At every layer, the products and technologies implemented were selected to maximize investment dollars and to ensure architectural integrity (i.e., Product A works with Product B). This architecture stack is currently used to deliver information and services to every major user community in State government.

Specific benefits of the architecture include:

- *Reduced costs for new applications*

- *Improved access to legacy data*

- *Centralized help desk, backup and recovery services*

- *Faster delivery of applications across a multitude of devices and networks*

- *Minimized data redundancy through data sharing*

- *Reduced dependency on proprietary components*

- *Reduced risk in reliable operations, security and change management*

- *Expert staff specially trained on enterprise platforms*

NJ IT Architecture Stack

24x7 Enterprise Help Desk	24x7 Enterprise Systems Management	Presentation & Portal Services	Identity Management	Performance Assessment
		Integration & Messaging		
		Application Development & Infrastructure		
		Data Management		
		Enterprise Servers and Operating Systems		
		Next Generation Services Network		
		Facilities and Environmentals		

While the existing infrastructure is described by way of an architectural stack, the State has undertaken an Enterprise Architecture program to focus on the Business, Information and Technology needs of the State as an enterprise. This program helped to achieve success in the government-to-business domain undertaken to promote the Governor's initiatives to stimulate economic growth and improve incentives for small and minority businesses. This initiative brought together executives from 21 State agencies to define the common vision for accomplishing this mission. This cross-agency cooperation will be used as a model to achieve success in growing the enterprise to satisfy requirements in other domains.

Facilities and Environmentals

NJ IT Architecture Stack

24x7 Enterprise Help Desk	24x7 Enterprise Systems Management		Identity Management	Performance Assessment
		Presentation & Portal Services		
		Integration & Messaging		
		Application Development & Infrastructure		
		Data Management		
		Enterprise Servers and Operating Systems		
		Garden State Network		
		Facilities and Environmentals		
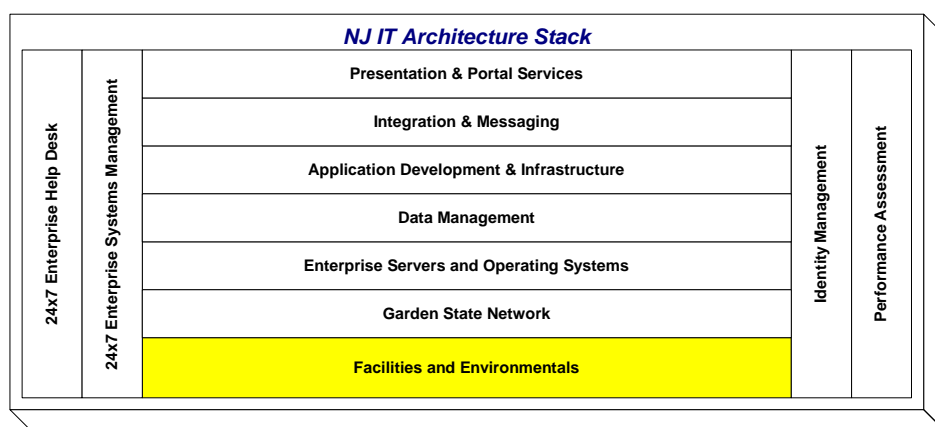
The State data center facilities maintain a symmetrical design in that the key infrastructure, system, and networking technologies have been duplicated between facilities. This common symmetry allows each facility to operate independently while providing back up services for its counterpart. High-speed fiber links both facilities allowing clients to freely deploy servers at the facilities. The data center facilities offer 24x7 complete operational and production services.

One of the facilities serves a dual purpose; that data center provides backup and recovery services for the mainframe environments and critical infrastructure services and serves as the State's disaster recovery facility. It will also be leveraged as a production hosting environment to support applications in conjunction with the other data centers. Plans are also underway to provide agencies with alternative geographic locations where mission critical applications can be hosted in the event of a disaster scenario at the primary facilities or as a means to provide additional capacity.

Control Center

Operation of the data centers are managed by a central Control Center. This control center is manned by a highly trained group of support professionals twenty-four hours a day, three hundred and sixty-five days a year. The responsibility of Control Center personnel is to ensure the availability, reliability and operational status of all production servers, the network, the environmental systems, and security systems within the facility. Facility Management, Capacity/Performance and Network Management systems and software are utilized by Control Center personnel to proactively monitor and display the status of these systems within the facility.

Next Generation Services Network (formerly Garden State Network)

NJ IT Architecture Stack

24x7 Enterprise Help Desk	24x7 Enterprise Systems Management	Presentation & Portal Services	Identity Management	Performance Assessment
		Integration & Messaging		
		Application Development & Infrastructure		
		Data Management		
		Enterprise Servers and Operating Systems		
		Next Generation Services Network		
		Facilities and Environmentals		
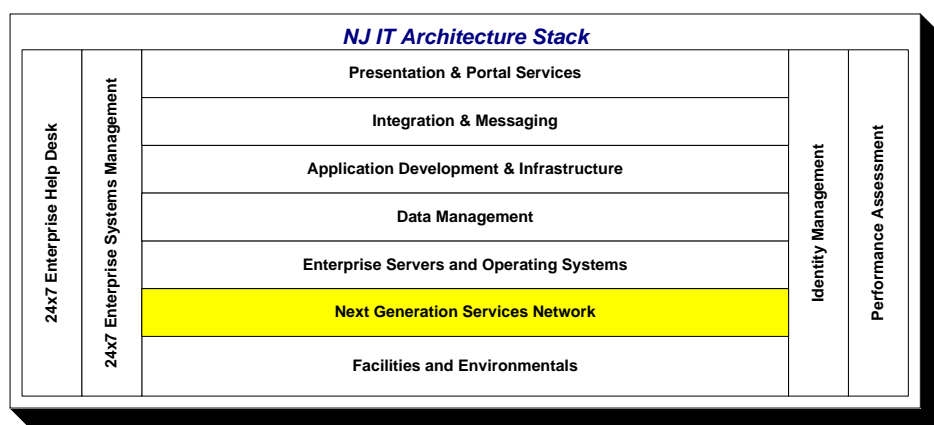
Network Services Architecture

It is the mission of the New Jersey Office of Information Technologies, Network Services Directorate to transform the Garden State Network and supporting infrastructure(s) into a Highly Scalable, Converged Network capable of fully supporting the Business, Public Safety and Public Access requirements of the Executive Branch Departments and Agencies, and Municipal First Responders and Law Enforcement.

The New Jersey Office of Information Technology is responsible for the Architectural Design, Engineering and Operational Control of the State of New Jersey's Garden State Network (GSN) and its ongoing transformation into the Next Generation Services Network (NGSN). Enterprise Network Services provided to State of New Jersey Departmental Partner and Client Entities are Internet Services, WAN Access, Aggregation and Backbone Services, Network Infrastructure Security and Access Control, Data Center Infrastructure, E-Commerce application Hosting, Secured Remote Access and Extranet Services, IP based mainframe Access and Unified Communications, VOIP and Video services.

Wide Area Infrastructure Transition

The transition of the legacy Garden State Network into the Next Generation Services Network continues. Carrier based technologies providing Cross LATA backbone services have been transitioned over to State of New Jersey owned fiber assets. Cross LATA backbone services, remaining legacy ATM and Carrier Ethernet are now transported via State of New Jersey managed MPLS (Multiprotocol Label Switching) services. The number of node facilities has expanded from six (legacy GSN) to a total of eleven, each supporting SONJ optical and MPLS transport services. All node facilities provide aggregation services for remote Departmental and Agency traffic and to facilitate NNI, Network to Network Interfaces, supporting the remnants of the ATM core infrastructure and Carrier Ethernet aggregation. Primary transport technologies servicing the Next Generation Services backbone are Optical Wave, CWDM, DWDM, Ethernet and MPLS. Departmental and Agency facilities (CPE's) connect either through their central nodes or to NJOIT NGSN node facilities (PE's) primarily with T-1, ATM, FRASI or Carrier Ethernet. The Inter-LATA traffic aggregation is supported via Verizon OC3, T3, OC12 and Ethernet Virtual Private Line (EVPL or Carrier Ethernet).

The Next Generation Services Network Architecture

The impetus for the development of the Next Generation Services Network was to capitalize on the potential synergies available through the use of governmental shared assets, leveraging available State of new Jersey owned infrastructure to develop a standard State-wide enterprise model for providing essential networking services, support for industry standard technologies such as 1 and 10 Gigabit Ethernet, FCOE, end-to-end Quality of Service to support IP based VOIP/Telephony, and Video Conferencing initiatives. Through the utilization of State of New Jersey owned dark fiber assets, the vision of building a State-wide fiber based network with protected on-ring presence for each of the State of New Jersey's data centers and core network locations has been realized and

fulfilled. With two of the major ring components completed (the Southern and Northern Rings and the targeted completion of the Central Dark Fiber Ring scheduled for Fiscal Year 2014), the NGSN is now positioned to add significant value to the State of New Jersey by supplying carrier transport and converged IP services to support Executive Branch operations, Public Safety initiatives and critical strategic objectives set forth by the Office of the Governor.

NGSN Strategic Benefits

The Next Generation Services Network provides expanded, on ring points of presence to deliver network access and backbone services to the State of New Jersey. The NGSN is comprised of 11 main node facilities located in each of the communication LATAs within the State. The NGSN provides a significant increase in bandwidth capacity and support for IP based services in comparison to the legacy Garden State Network. The NGSN core, interconnecting the State's data center facilities, currently provides 20 gigabits capacity on a protected fiber ring. The completed and operational outer ring components, the Southern and Northern rings currently provide 8 gigabits of backbone capacity, bi-directional with a near-term target capacity of 20 gigabits on protected fiber. The NGSN supports full convergence of networking services, data, voice and video, end-to-end Quality of Service (QOS) and Private Virtual Cloud Services.

NGSN Ongoing Migration Strategy utilizing Carrier Ethernet Services

OIT has co-located the NGSN optical ring with the Carrier Ethernet Networks in each communication LATA and are using OIT managed MPLS technologies to seamlessly provision the Carrier Ethernet Services for our State of New Jersey Departmental and Agency IT partners. The utilization of Carrier Ethernet to support remote client facilities enables the State's MPLS Services to provide path isolation through the use of L2 and L3 virtualization to support isolation of Departmental and Agency traffic and to rapidly provision bandwidth to support increasing capacity and IP services demands. Multi-tenant facilities supported through Carrier Ethernet Services enable OIT to implement a shared services model utilizing a protected Ethernet circuit, OIT managed router and switch, multiple sub-interface configurations to support traffic isolation and individual tenant capacity demands over the MPLS enabled backbone.

Ongoing Legacy GSN Migration Strategy

State of New Jersey Departmental and Agency partners not yet positioned to migrate to MPLS supported Carrier Ethernet Services require OIT to deploy Carrier ATM OC at NGSN Nodes to support aggregation of legacy ILEC Frame Relay (FRASI, Frame to ATM) and ATM network to network (NNI) circuits. OIT provider edge routers will provision the Cross-LATA ATM links over the State's MPLS backbone. This strategy is currently implemented in the Southern and Northern LATA's enabling OIT to migrate off a major section of the legacy ATM core. **NJOIT's target date to complete the full migration from legacy ATM and FRASI to Carrier Ethernet Services is the 3rd quarter of 2015.**

Advanced Services Supported on the Next Generation Services Network

- Real-time Application Services, Voice and Video Applications
- Voice Gateway/Dial Tone/Call Manager/Call Recording
- Video and Teleconferencing network isolation
- Secure Guest Access
- Wireless Network Isolation
- Wireless Guest Access
- Robust Data Center interconnects

Internet Services

Current Internet Services (circuits) are contractually two (2) 1 Gigabit Ethernet circuits supporting the State of New Jersey's Ingress/Egress Point of Presence at two physically diverse NGSN node facilities.

E-Commerce (N-Tier) Architectures

The State of NJ supports multiple multi-tiered environments in which to host E-commerce applications. A Cisco Catalyst 65xx fabric which provides network redundancy and a recently implemented Cisco Nexus 7K fabric which provides redundancy and High Availability services.

The multiple n-tier environments provide secure, but direct access to the State of New Jersey informational and critical line of business application systems. Current security policy dictates that web access directly from the Public Internet is limited. The n-tier environment provides presentation, business logic and data layers. Recent enhancements to the State's data center E-commerce/data center hosting environments include:

- Redundancy at all Network layers
- Redundant network connections for all servers
- Redundant power grids
- Increased throughput
- Access Policy Enforcement
- Integrated firewall service modules in Fail-Over Configuration
- Multiple Security Zones support 2-tier, and 3-tier
- Protection and Prevention Systems, Monitoring and Logging
- Network Services Distribution Model
- Simplified Cable Management for servers, SAN, KVM, IP DRAC

NGSN (N-Tier)

- The new application and network services hosting infrastructure will introduce a number of new features and services that the New Jersey Office of Information Technology can provide to Departmental and Agency partners. The Nexus architecture will provide a comprehensive suite of services to host customer servers and applications, such as load-balancing, global site selection, security as a service and global redundancy, cross site high availability services, private cloud and virtual multi-tenancy hosting services.

Virtual Datacenter Architecture

- Allows seamless migrations of application/virtual hosts between data centers, Overlay Transport Virtualization Services (OTV, L2 encapsulation)
- Work Load Mobility and Load Balancing
- Consolidation and Migration
- Maintenance enablement
- Disaster avoidance (with Storage VMotion or active/active replication storage models)
- Future Architecture
- Private Cloud, Inter-Cloud, Multi-tenancy

Disaster Recovery, High Availability, Inter-Intra Datacenter

- OTV, WDM Optical Site Interlinks
- Cross Site virtualation, application clustering, Microsoft Exchange, etc
- Cross Site application load balancing, application availability probing, Supports Active/Active, Active/Passive

Architecture Design Benefits

- Modular topology—building blocks provide flexibility during changes and upgrades

- Effective virtualization capabilities, building blocks for Cloud Computing capabilities
- Creates small fault domains—clear demarcations and isolation
- Incorporates balance of Layer 2 and Layer 3 foundation, leveraging the strength of both
- Allows integration of load balancing and security
- Promotes deterministic traffic patterns
- Easier to grow, understand, and troubleshoot
- Highly flexible design

  ➢ Infrastructure standardization
  ➢ Address OOB objectives
  ➢ Secure segmentation

- Support different Application Redundancy Models including Active/Active Front End and Active/Active Back End

Secure Remote User Access

The State maintains several mechanisms to provide secure remote user access to internally hosted infrastructure and application resources:

- The State of NJ Enterprise Portal, which provides access for thousands of users to core computing resources via HTTP Proxy services and a proprietary application VPN service (see State Portal Overview). For applications that do not meet the traditional E-commerce model for web, presentation and data layer design, extranet connectivity is available.

- Extranet connections require point-to-point or MPLS enabled connections from the extranet partner to the extranet firewall infrastructure either via a point to point data circuit, or through an IPSec tunnel across the internet. The cost of these connections varies based on the type of data circuit ordered, and the equipment required to terminate the circuits.

- The Remote Access VPN solution provides SSL and IPSEC VPN services to State and non-State users. VPN services are only available to system administrators to provide off-site access for system maintenance and monitoring. State employees are required to register for two factor authentication to the VPN. Non-State users (i.e., consultants) are required to register through the State agency for two factor authentication to the VPN.

- GOTOMYPC services are available to limited State employees. GOTOMYPC services provide for business continuity in order to access computer systems from home when a facility is inaccessible. State employees are required to register for two factor authentication to the GOTOMYPC services.

- Citrix services are available for both State and non-State users (consultants). Application development processes can utilize Citrix services for off-site access to maintain code enhancements and conduct application testing. State employees are required to register for two factor authentication to the Citrix services. Non-State users (i.e., consultants) are required to register through the State agency for two factor authentication to Citrix services.

- Air cards use for remote communications are being implemented which provide access to the State's GSN via the Remote Access VPN solution and the Internet directly. Security policies are defined and established to limited use, and protection and prevention.

- Dialup services are being phased out in favor of higher bandwidth service options. No new dialup users will be accepted at this time.

Network Service Level Management

Real-time proactive monitoring is performed on all OIT managed network infrastructure devices, both Wide Area Networks (GSN, NGSN) and the data center core and E-commerce hosting environments. All devices are monitored for performance, availability and health statistics, including CPU and memory usage, operational temperature control, fans, power supply, individual interface and sub-interface events, Routing Protocol status and other hardware and software event identification. Event Management and correlation with root cause analysis are used to assist in identifying and resolving critical problems. This NSM intelligence enables OIT network operations to pinpoint the underlying cause and overall impact of critical network infrastructure events thus reducing down-time and maximizing infrastructure availability.

Tivoli NMIP (NetCool) Service Portal
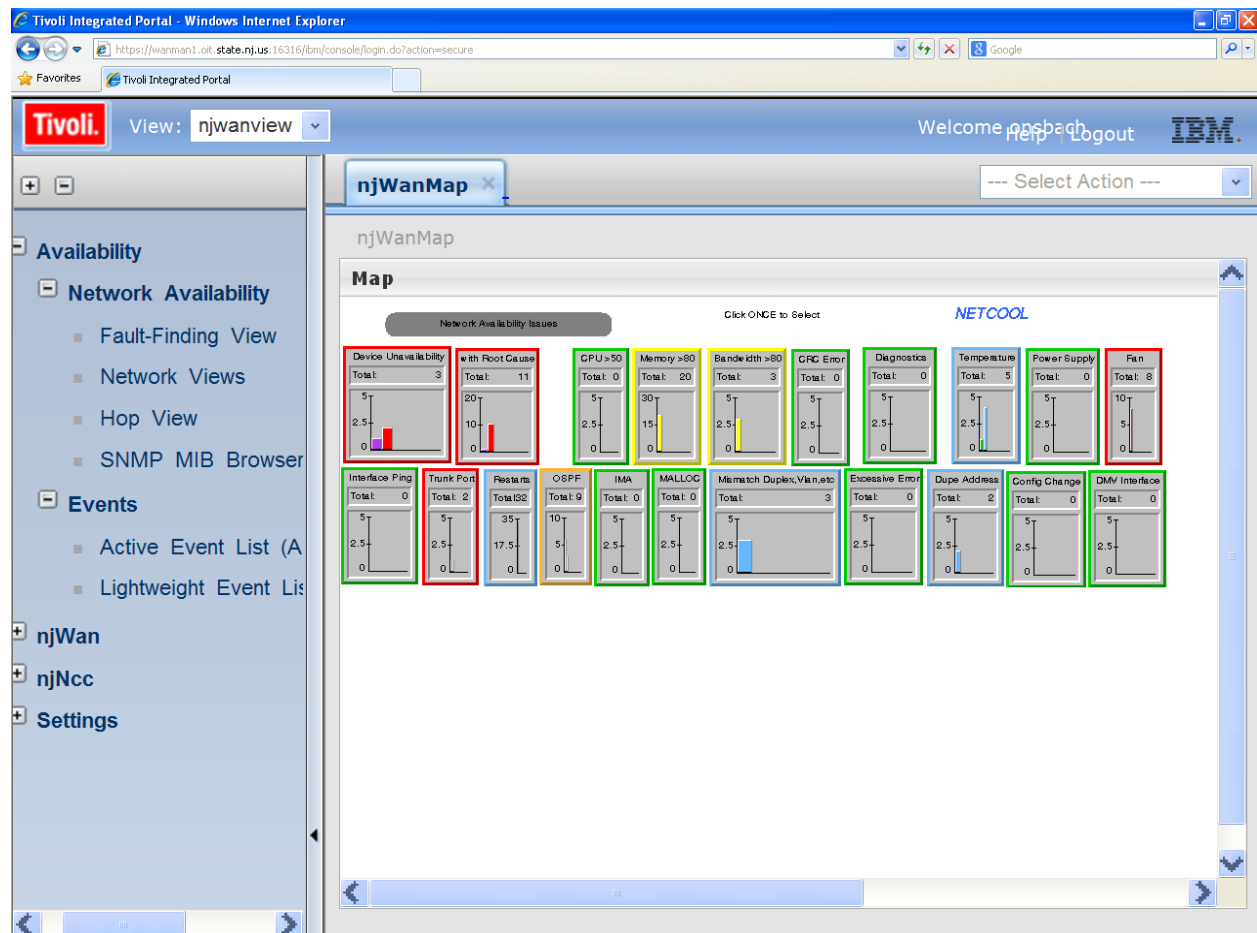
Network Performance Assessments, Application Triage and Performance Service Level Monitoring

The New Jersey Office of Information technology currently utilizes multiple product sets to Monitor, Assess, Diagnose and provide Fault Domain or Root Cause Analysis for Network and Application Performance issues. Software and hardware based network and application probes are deployed to perform baseline analysis of existing network environments prior to deploying new applications, upgrading existing applications or implementing new IP services such as VOIP or telephony applications. Application protocols, their respective traffic volumes traversing the local (LAN) and wide area network (WAN) are identified and their bandwidth consumption, average response

times, RTT's Round Trip Times (TCP hand shake), conversation pairs and traffic volumes measured. This analysis can be used as a benchmark comparison against future application or network performance.

Application Pre-deployment assessment services are available for Departmental and Agency Clients to assess the network performance characteristics of individual application functions or transactions. Characteristics such as TCP windowing, packet size, conversation flow, node sending and processing behaviors and inspection of processing threads between clients and servers or between nodes in a N-tier hosting environment are examined. The assessment process also includes performance modeling based upon variations in available bandwidth, processing performance and modification in TCP/IP stack parameters.

In the production environment, Application Triage tools are deployed to monitor and diagnose degradations in application performance and to determine the root cause(s) of poor application performance. Triage tools help to determine whether poor application performance are the result of underpowered hosts, the network infrastructure, the application code or an inefficient host server platform/OS or database.

Pro-active/Real-time Monitoring of Application Service Levels is the newest offering available to Departmental and Agency Clients who have deployed application systems within the N-tier environment at the State's data centers. Deployed Performance SLM toolsets monitor end to end flow of application traffic traversing the E-Commerce environment. The SLM process provides both real-time and historical information regarding the performance, usage and availability of key business applications. The SLM toolsets provide deep dive analysis capabilities of Web based application components including SSL decryption, HTTP page load sequencing, hit level and error analysis, application protocol analysis and decode, including database performance metrics, tiered fault domain isolation analysis, available in real-time or through an extensive set of data mining services.

Mainframe Access

As of Fiscal Year 2012, all legacy System Network Architecture (SNA) (external to physical mainframe) infrastructure, SDLC circuits, controllers, SNI gateways and other supporting infrastructure have been eliminated from the OIT managed network infrastructure. Standard industry protocols are in use to access the mainframe.

Performance monitoring and PD tool in use is CA's NetMaster. VTAM routing tool/product is William Data Systems RouteView, used to define network paths. Mainframe support product is IBM's NetView, used to monitor network components and allow Network Call Center/System Command Center access for the mainframe network.

Enterprise Servers & Operating Systems

		NJ IT Architecture Stack		
		Presentation & Portal Services		
		Integration & Messaging		
		Application Development & Infrastructure		
24x7 Enterprise Help Desk	24x7 Enterprise Systems Management	Data Management	Identity Management	Performance Assessment
		Enterprise Servers and Operating Systems		
		Next Generation Services Network		
		Facilities and Environmentals		
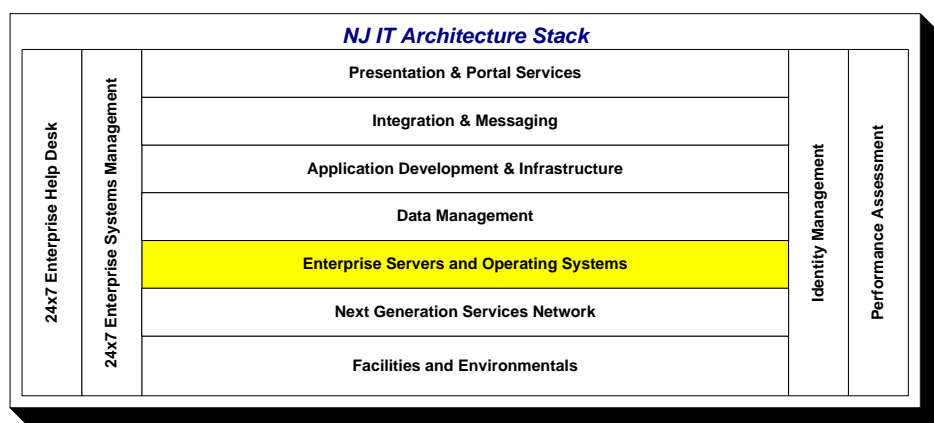
Shared Server Infrastructure

Mainframes and servers are centralized to offer a common location to manage the distributed environment. Cabinets are provided to rack servers and eliminate excess footprint. Implementation of a standard KVM (Keyboard, Video, Mouse) matrix switching backbone solution at both facilities has improved floor space utilization, cable management and server access as well as reduced equipment requirements and power consumption. Optimizing key server resources through common logical and physical environments positions the State to properly plan, manage and control a growing server infrastructure. For all servers housed in this environment, OIT and the agency may share the administration of the solution components.

Based on the best-supported environments by the IT community, the SSI supports multiple operating system platforms.

Server Virtualization and Consolidation

Another key data center optimization strategy pursued by OIT is server virtualization and consolidation. Implementing this strategy is dependent on technological advances in both hardware and software that have now cut across all operating system platforms noted above. This approach saves on data center floor space, power, and cooling per unit of processing capacity. In addition, operations, administration, and maintenance can be addressed more efficiently and less expensively. Consequently, for new applications, OIT is driving the deployment of virtualized servers as the preferred approach. For existing applications, OIT is pursuing server virtualization and consolidation where it makes business sense to do so (e.g., at the point of equipment refresh or maintenance renewal).

OIT is also pursuing virtualization and consolidation of infrastructure services as more and more agencies leverage the State's enterprise hosting architecture. Two specific examples are given below.

Shared Hosting Services

OIT redesigned the architecture that provides for enterprise clustered Citrix farm authentication via the Windows AD that enables agencies to leverage an enterprise hosting environment. Application development processes can also use Citrix services for off-site access to maintain code enhancements and conduct application testing.

Enterprise Internet Gateway

State agencies currently use multiple product sets to monitor, access, report, track, and manage Internet access. OIT has completed the build and implementation of an enterprise architecture to provide these services to address security for inbound and outbound risks at the lowest total cost of ownership. This environment provides:

- A unified and central management analytic system that lets each agency leverage the power of dynamic web communication, centralized database, central State governance and prevention of web application attacks
- A unified management console that gives cleaner visibility into what is going on statewide with web, e-mail, and data security providing each agency the ability to set their own rules and policies without affecting others
- Enforcement of the State's Internet use policies
- Cloud or hybrid functionality for State and contract employees without putting the State network at risk and the power of on-premise and off- premise Internet access with the flexibility of the cloud, reducing complexity, increasing effectiveness and lowering overall costs
- A single point of entry for all web access in and out of the State enterprise network systems to help mitigate vulnerability to State systems
- A redundant implementation to allow all agencies to fail over should their primary appliance fail
- A central archive for discovery and audit reporting on Internet access violations and compliance

Enterprise Software Distribution

OIT recently completed the build of an enterprise Endpoint Management solution suite. This suite offers a software family suite of products. The product can help reduce the cost of deploying and managing desktops, notebooks, and thin clients from a centralized location in your environment, improve efficiencies with comprehensive configuration management, provide rapid and reliable Windows planning, migration and ongoing management of systems, take control and automate IT infrastructure, and much more.

Storage Area Network

The State manages a Storage Area Network (SAN), Storage Management offers fully redundant storage arrays, with over 4.5 PB of storage currently in use. The SAN consists of a redundant core to edge fibre channel communication

that provides physical connections, a management layer that organizes the connections and storage layer that controls data delivery and security. Storage devices are connected to servers in a networked fashion, using directors to build the topology.   The State uses a variety of storage array types to optimize performance and minimize price based on storage needs.

The SAN currently supports connection speeds of 2, 4 and 8 GB.

In order for a server to "talk" to the SAN, an additional piece of hardware called a Host Bus Adapter (HBA) must be installed in the server. Two HBAs are needed in order to provide redundant paths to the SAN; this eliminates the possibility of having a single point of failure. Once connected, disk space can be allocated from the storage array(s) and dedicated to a server. SAN technology presents many benefits to server data storage, such as:

- Centralized storage management
- Ability to add disk capacity dynamically
- Ability to replace a deficient server without loss of data
- Faster response time than internal SCSI disks
- Potential for improved backup and disaster recovery techniques
- Better storage attributes – hardware RAID, dynamic sparing, remote data copy, mirroring, and more.

Storage Management also offers boot from SAN. Using this method, all OS drives are replicated to the OARS recovery site for quicker server recovery.

Backup and Restore Services

OIT Storage Management is currently converting to NetBackup for backup and restore services available to clients within the multiple security zones.

Clients consist of Windows, Linux, Solaris, AIX, Novell and VMware systems, as well as Oracle, and SQL, databases. Other clients are available upon request.

These services require NetBackup software loaded on the target server that selects the data for backup on the server, and then sends the selected data to the NetBackup server by the way of TCP/IP.

Storage Management requires the creation of a User ID with Root/Administrator authority on the target server, which is used to install the client, monitor backups, and troubleshoot any problems that may occur during daily backup processing.

For all servers at the State's data center facilities, an additional Network Interface Card (NIC) should be installed and connected to the Storage Management Backup Network in order to reduce the backup window, and eliminate network contention.

Note that these services are for backup and restore of the server data only. Data archiving is a different process.

Basic Server Backup/Restore Policy (Unstructured Data)

The standard client backup begins at midnight 0000 hours (12:00 AM) with backup duration dependent upon client hardware and network bandwidth. Most clients are usually finished the backup processing by 0600 hours (6:00 AM), and must be completed by 0730 hours (7:30 AM).

The first backup is that of a full system, meaning that every file not specifically excluded by the NetBackup configurations files is sent to the backup server. Subsequent full backups are done every 8 weeks. Incremental backups will be done in between so that only those files that have changed since the last backup are sent to the backup server. This method reduces network bandwidth consumption and backup storage requirements. Every 4 weeks a Synthetic Full will be created. A Synthetic Full creates a new tape merging the full with all the incremental taken since the last full. Backup data is stored on virtual tape.

The standard backup policy will retain unstructured data for a period of 60 days.

Structured Data Services

NetBackup will fail to backup files that are open for writing. For data that must be available to an application 24/7, NetBackup provides other clients that must be utilized.

Oracle Database Backups

Storage Management uses the Oracle Recovery Manager (RMAN) in conjunction with NetBackup to backup and restore Oracle database instances. In most cases, clients depend on 24/7 operations that cannot be interrupted for backup processing.   Storage Management utilizes a hot backup procedure that allows database operations to continue while the database is backed up unless the client has specified otherwise.

The standard RMAN hot backup policy consists of a full backup of the entire database once per week. Backups are also performed on a nightly and non-cumulative incremental basis for the remainder of the week.

Control files are backed up nightly along with the full and incremental database backups. Archive Logs are also backed up via the nightly RMAN scripts unless the logs are managed by the migration client. Oracle parameter files, password files, and other configuration files are not managed by RMAN, and should be backed up using the NetBackup server.

All RMAN backups are tracked by an RMAN recovery catalog residing on the backup server, and all backup pieces generated by RMAN are stored by NetBackup. Storage Management offers a recovery window of twenty-one (21) days for the standard Oracle client.

This means that Storage Management keeps all RMAN backups necessary to restore a database to a point in time equal to twenty-one (21) days prior to the current time or today minus twenty-one (21) days. Once an RMAN backup piece is no longer useful for this recovery window, it is expired and no longer available for restore operations.

SQL Database Services

A client add-on for the NetBackup can be installed and configured on each server Microsoft SQL Server version 2000 or greater.

A full "hot" backup is done on Friday evening on each Netbackup client. The "hot" backup is an open-file-supported-backup of each active database residing on the Windows server.

Differential "hot" backups are performed from Saturday through Thursday. Incrementals include all files that have been changed since the last full backup. Log "hot" backups are also available, and may be set to occur at any frequency greater than twenty (20) minutes between nightly incremental backups.

The standard backup policy will retain up to twenty-one (21) backup versions of a file as long as that file exists on the server. Once the file has been backed up twenty-one (21) times, the oldest backup copy of the file will be expired with each subsequent backup.

Exchange Services

A client add-on for NetBackup can be installed and configured on each server running Microsoft Exchange Server.

A full "hot" backup is done on Sunday evening on each NetBackup client. The "hot" backup is an open-file-supported-backup of each active Exchange instance.

Incremental "hot" backups are performed from Saturday through Thursday. Incrementals include all files that have been changed since the last full backup.

The standard backup policy will retain up to twenty-one (21) backup versions of a file as long as that files exists on the server. Once the file has been backed up twenty-one (21) times, the oldest backup copy of the file will be expired with each subsequent backup.

Archiving Services

The Archiving Service is accomplished using Hitachi content platform. It allows a client through policy to manage data retention periods and data destruction.

Today's uses:
- State police video retention
- Oracle archive logs
- SQL archive logs
- User data

Archiving can be attained through 4 methods:
1. NFS mount
2. Cifs mount
3. Rest/Https protocol
4. 3rd party plugin

Methods 1 and 2 require the client to mount a networked drive to their system. This requires a network connection to the storage management backup network. Method 3 uses HCP utility which can be run from crontab in UNIX or task scheduler in windows. It connects to the archive and stores the requested data.

Data archives to either data centers and then is replicated to the other data center.

HDI Service

HDI – Hitachi Data Ingestor is useful to replace remote and local file servers. It can ingest existing file servers and then replace them. Once the HDI ingests the data it transports a copy to the data centers and is replicated to the

alternate site. Today this service is only available to one facility but plans are in the works to place one in another facility to take advantage of this technology.

Bare Machine Recovery Services

Bare machine recovery is the process of recovering a server instance to different hardware that may be similar or dissimilar in configuration. This service will be pursued in the near future.

Data Management
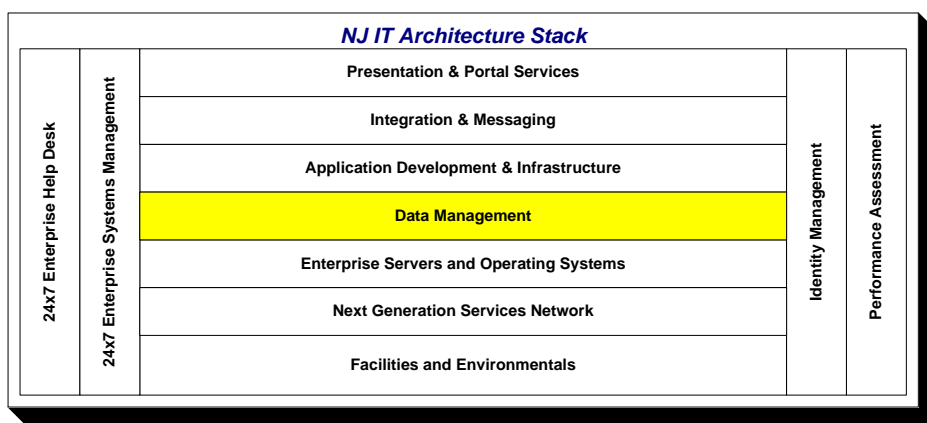
New Jersey's Enterprise Information Architecture, embodied as the New Jersey Data Governance Framework, is captured in the following components:

- The *New Jersey Data Governance Framework Strategic Plan* – This is the conceptual road map that describes the overarching New Jersey Enterprise Information Architecture.

- The *New Jersey Data Governance Framework Implementation Plan* – This is the guide for implementing these logical and physical components of the New Jersey Enterprise Information Architecture:

  o The *New Jersey Enterprise Information Management Framework* (**NJEIMF**) – This defines the logical information architecture, including use cases, design patterns, and the governance required.

  o The *New Jersey Reusable Data Infrastructure* (**NJRDI**) – This documents the physical solutions data stores embodied by the data management and integration technologies and methodologies.

  o The *New Jersey Enterprise Data Resource Inventory* (**NJEDRI**) – This documents the data architecture for the enterprise; the standards, common data structures, and reusable reference and master data of which all systems benefit to minimize costs, reduce time to deploy, and improve efficiencies.

To implement its Data Governance strategy, the State has created the New Jersey Enterprise Information Management Framework (NJEIMF), the New Jersey Reusable Data Infrastructure (NJRDI), and the New Jersey Enterprise Data Resource Inventory (NJEDRI) to deliver enterprise data management to the State's executive branch departments and agencies.

The NJEIMF is the enterprise information architecture for New Jersey - the art of expressing a model or concept of information used by complex or inter-related technology systems. It is a set of rules that determine what, and how and where, information will be collected, stored, processed, transmitted, presented, and used. The NJEIMF is summarized here in the New Jersey Shared IT Architecture document. It is also a separate document available on the NJ.gov web site.

The NJRDI is the data management infrastructure for New Jersey. The data management domain encompasses the collection, definition, and maintenance of data as well as the use and presentation of information derived from that data. The NJRDI provides the common tools and methodologies for defining data and implementing data management solutions consistent with the NJEIMF. The components of the NJRDI are documented in Appendix A of the New Jersey Shared IT Architecture document. This forms the basis for New Jersey's data management service delivery.

The NJEDRI includes the enterprise logical data model for New Jersey which describes logically data of interest to all or part of an entire enterprise. The NJEDRI also defines and standardizes data used to conduct business operations across business units. Finally, it documents practices and processes that optimize data management in the State. This document is available by request from the NJ OIT Data Governance Office.

Information Architecture

Information Architecture standardizes the design, definition, and relationships of the State's data elements, provides for the management of those data elements, and guides the creation, maintenance, and availability of the data. Its goal is to make data reusable to the greatest extent possible while improving overall data quality.

Consistent with the version 2 of the Data Management Association (DAMA) Data Management Book of Knowledge (DMBOK), data governance is the common driver for all of the NJSDI information architecture components. Data governance interacts with multiple touch points within the infrastructure, as described below.

Data Architecture

This is the overall relationship of data and data structures within the State of New Jersey. It provides an enterprise viewpoint to the more physical aspects of the information architecture.

Data Modeling & Design

This creates logical and physical definitions of data objects, providing for well-defined non-redundant logical structures that form the basis of all physical database implementations.

Data Storage & Operations

This manages the life cycle of the data asset at rest, typically within a relational database management system. It includes tiered capabilities to meet the storage requirements of different categories of data. It also includes backup, recovery, and restoration capabilities.

Data Security

This manages the confidentiality, integrity, and accessibility of data assets so they are available when and where needed but protected from unauthorized access or alteration.

Data Collection

This is provided by application development, through acquisition of commercial-off-the-shelf (COTS) software, and by importation of data from external partners and systems.

Data Storage

This manages the life cycle of the data asset at rest, typically, but not always, within a relational database management system. It includes tiered capabilities to meet the storage requirements of different categories of data. It also includes backup, recovery, and restoration capabilities.

Data Integration & Interoperability

This manages the delivery and receipt of data in motion. Integration brings together and rationalizes data from two or more systems to create an enhanced data asset not otherwise provided by any one system. Interoperability enables otherwise independent systems to operate as a single virtual system.

Documents & Content

This is the world of unstructured data that must be integrated with structured data. Not all unstructured data fall into this category. For the data that does, this domain provides technologies and practices to make unstructured data consumable in a traditional data management environment.

Reference & Master Data

Reference data are the relatively static data sets, typically in the form of codes and descriptions which categorize other data throughout New Jersey. Master data domains represent the descriptions of people, places, and things that change slowly and have core attributes that are used throughout New Jersey.

Data Warehousing & Business Intelligence

This is the repurposing of data captured in transactional systems for analytic and decision-support purposes. It is the delivery of information to different user communities based upon their individual requirements, using graphical end-user tools. The data is formatted as much as possible to anticipate reporting needs, and may be presented differently to different groups, but always from a common source for consistency.

Metadata

This documents logical and physical definitions of data objects and their relationships to each other. This documentation is consumed by the other information architecture processes and makes them more effective.

Data Quality

These are the processes to identify, address, improve and manage the overall integrity, accuracy, and usefulness of the enterprise data assets.

Data Governance

Data governance is a set of processes that ensures that important data assets are formally managed throughout an enterprise. Data governance ensures that data is defined, has a known level of quality, and can be used for the intended purpose; in other words, it can be trusted.

New Jersey data governance is focused on identifying those individuals and organizations with the role of defining data objects, identifying the authoritative source for each data object, and classifying each data object. It assists in the resolution of data quality issues, so that New Jersey state government can become more efficient.

Data Steward

The Data Steward is the individual or unit that manages the authoritative source for a particular piece of data and controls its definition and access. A Data Steward is not the same as a Data Custodian, an individual or unit that has been assigned the duty to manage the data under the direction of the Data Steward. A Data Steward is not the same as a Data Owner, which can be a third-party person or organization that the data describes and that has provided the data to the State when requested or required by a State agency.

Data Tiers

New Jersey categorizes data into four tiers – Universal, Enterprise, Line-of-Business, and Programmatic. These data tiers provide a way of framing data governance and data steward responsibilities as well as helping to define the scope of data modeling and data management efforts.

- **Universal (Tier 0)** refers to data commonly referred to as Master Data. This is data that describes persons, places, or things independent of their relationship with the State.

- **Enterprise (Tier 1)** refers to data that is common across all State agencies but within the context of their own organization, such as Financial, Asset, and Human Resources data.

- **Line-of-Business (Tier 2)** refers to data that is common across a particular line-of-business involving more than one agency, such as social services data, business community data, or early childhood data.

- **Programmatic (Tier 3)** refers to data that is specific to a single program area within a single agency and is unlikely to have value outside of that context.

NJ data is categorized into four integration tiers:
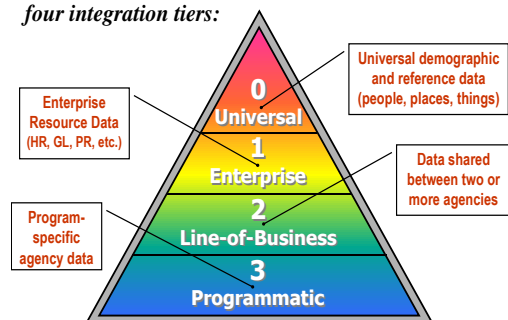
Information Asset Classification

The State has implemented an Information Asset Classification policy to address enterprise security for information assets and data management. Information classification is the categorization of data for its most secure, effective and efficient use. Classification assigns data a level of sensitivity, criticality, and/or potential loss impact as it is being created, amended, enhanced, stored, or transmitted. Classification of the data will also determine the extent to which the asset needs to be controlled or secured and is also indicative of its value in terms of Business Assets.

New Jersey requires that all data maintained by the State be classified as to its Confidentiality, Availability, and Integrity risk, in accordance with the FIPS 199 standard.

- Confidentiality – The need to preserve authorized restrictions on information access and disclosure, including the need for protecting personal privacy and proprietary information.

- Integrity – The need to guard against improper information modification or destruction, including ensuring the non-repudiation and authenticity of the information.

- Availability – The need to ensure timely and reliable access to and use of information.

New Jersey's Model-Driven Development (MDD) Approach

New Jersey's information architecture requires a model-driven approach to development. Where it exists, this process begins with an existing Conceptual Business Model (CBM). The CBM guides the development of a Solution Conceptual Data Model (CDM) to capture the key information needs of the business. This model is created with participation by stakeholders at the highest levels of the business. The CDM guides subsequent modeling efforts and documents an overall view of the business, even for areas outside the scope of the application being developed. The CDM feeds back into the organization's CBM, or forms the basis for one where it does not yet exist.

After creation of the CDM, the logical modeling process captures detailed user requirements and business rules. A Solution Logical Data Model (LDM) is created representing the scope of the project. It is consistent with both the CDM, the Logical Business Model (LBM) for the business or subject area where one exists, and with the relevant components of the NJ Enterprise Data Model (NJEDM). The LDM is normalized, fully attributed, and consistent with the NJ Data Naming Convention. The entities and attributes of the LDM are registered with the NJ Data Architecture unit and recorded in its Data Registry. The NJ Data Architecture unit validates and approves all entity and attribute names. The LDM feeds back into the organization's LBM and the NJEDM.

Once the LDM is created, it is used to produce a Physical Data Model (PDM) for the project. It is in this PDM that any changes to data structures to address performance, security, or development issues are made. The LDM remains fully normalized and the physical changes are mapped from it. Once the PDM is approved, it is used to generate the Data Description Language (DDL) needed to create the actual database structures required.

NJ Data Architecture Workflow
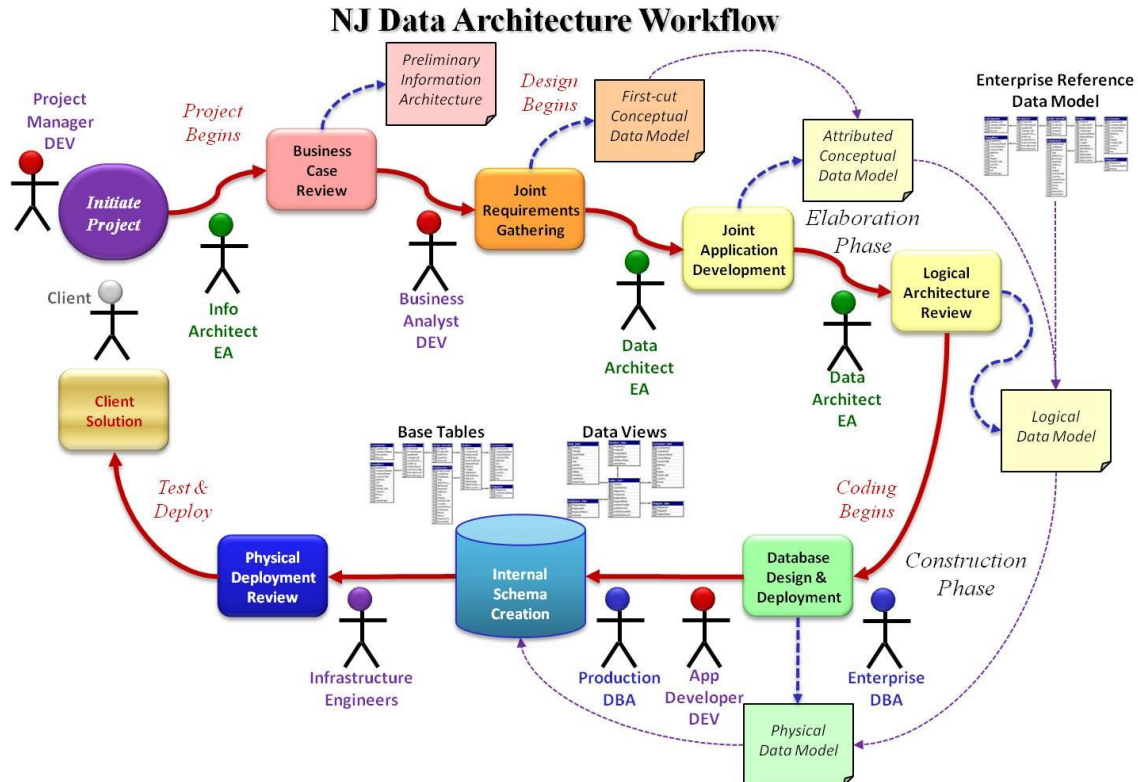
Whenever possible, changes that need to be made to the application after the initial database is created should first be made in the LDM. The changes are then progressed through the PDM to the actual database. In this way, the

documentation remains accurate and synchronized, and the impact of changes on data integrity is fully understood. In cases where changes must be made immediately to the physical database to correct an urgent production problem, it is imperative (and required) that the developers update the LDM and PDM immediately thereafter.

National Information Exchange Model (NIEM)

The National Information Exchange Model (NIEM) is a national XML-based information exchange framework. NIEM represents a collaborative partnership of agencies and organizations across all levels of government (federal, state, tribal, and local) and with private industry. NIEM is designed to develop, disseminate, and support enterprise-wide information exchange standards and processes that will enable jurisdictions to automate information sharing.

NIEM is not a software program, database, network, or computer system. NIEM facilitates the creation of automated enterprise-wide information exchanges which can be uniformly developed, centrally maintained, quickly identified and discovered, and efficiently reused. As a data model NIEM is a hybrid of multiple model types:

- It is a logical reference data model, in that it documents the business definitions of data of interest to multiple organizations and jurisdictions.

- It is a collection of logical business models, in that it documents the entities and attributes of multiple subject areas or domains.

- It is a physical data model, in that it provides XML schemata that can be used to exchange physical data, and these schemata are mapped back to logical data definitions.

New Jersey Information Architecture Design Patterns

Design patterns are a way to implement – to apply – information architecture. The term "design pattern" is used to refer to any formal, documented approach which addresses issues of software architecture, design, or programming implementation. A design pattern provides a scheme for refining the subsystems or components of a software system, or the relationships between them. It describes a commonly recurring structure of communicating components that solves a general design problem within a particular context.

A design pattern provides a formal definition of a solution and of the problems to which it applies. The goal of design patterns is to avoid approaching each situation as a problem that has never been seen before and, instead, to make it possible to repeat solutions that have worked. In particular, a design pattern distills the experience of an expert or the best practices of a community so everyone can apply that expertise. While the approach originated in architecture and has seen great success in software engineering, design patterns apply equally well to information architecture.

Information architecture design patterns can be organized into several broad categories. A particular solution may combine two or more design patterns. It is important that a single design pattern not be extended to meet needs that are better met with a different design pattern. This is a common mistake which introduces true complexity into the enterprise with the false hope of reduced complexity.

Data Collection

A Data Collection Design Pattern is for the processing of transactions and events, either as stand-alone applications, or with real-time or batch consumption of data from other transactional or reference data systems. Such systems are responsible for ensuring the integrity of the data transactions that they process.

Examples of Data Collection Design Patterns include:
- Traditional mainframe batch processing system
- Distributed N-tier transactional system

Data Reporting

A Data Reporting Design Pattern is for the reporting, analysis, or publication of integrated, aggregated and summarize data. Such systems are responsible for delivering data in the format required to those who need it as efficiently as possible.

Examples of Data Reporting Design Patterns include:
- Operational reporting
- Analytical reporting

- Data mining

Data Integration

A Data Integration Design Pattern is for the aggregation and integration transactional and reference data from multiple systems within an enterprise to provide current or historical integrated data as a source for reporting, analysis, or system integration, or, for the management of common reference or master data used by multiple systems.

Examples of Data Integration Design Patterns include:
- Operational Data Store
- Enterprise Data Warehouse
- Master Data Management through a MDM platform
- Real-time service bus-based data inquiry through a common exchange model

Data Exchange

A Data Exchange Design Pattern is for the exchange of data with an external partner. Such systems are responsible for real-time inquiry and response or for traditional sharing of data (file transfer). Whenever possible, these exchanges should be based upon a common governance model.

Examples of Data Exchange Design Patterns include:
- Real-time web-service-based data inquiry through a common exchange model
- Real-time web-service-based data inquiry through a negotiated exchange model
- Data sharing using file transfer with a negotiated exchange model

New Jersey Design Pattern Principles

Transactional System to Collect Data

To the greatest extent possible, new transactional system physical designs shall be developed using a fully normalized logical data model consistent with the NJERDM and the State's naming convention. These systems shall be hosted within an industry-standard SQL-enabled relational database management system (RDBMS), and shall use to the greatest extent possible the referential integrity and domain constraint capabilities of the RDBMS to enforce business rules. These systems shall subscribe or consume common reference and master data defined and provided at the enterprise level.

Batch Integration of Inbound Data

Previous assumptions that batch processing windows will always be available to handle any size batch processing requirements are no longer valid. New batch processes must determine if processing smaller batches more often (even in near real-time as batches of one), processing batches while the systems are online, partitioning data or systems, or creating parallel processes are appropriate to achieve the goal of the process.

Real-time Integration of Data

Where there is a need for real-time integration of data, it shall be implemented as a web service. The format for real-time integration shall be defined in XML consistent with the NJERDM. Where one exists, the enterprise service bus (ESB) shall be used.

Provide Data to External Systems from Mainframe Systems

Because data used by one system may be of value to others, and because of the costs associated with creating multiple interfaces on mainframe systems, and because of the complexity of managing outbound interfaces in a mainframe environment, point-to-point solutions shall not be created. Instead, data required by an external system that is not already in the Enterprise Data Warehouse (EDW) environment shall be output to the EDW. The external system will either pull or have pushed to it the data from the EDW.

Internal Reporting of Operational Data

Complex reporting needs should not be processed in real-time against critical or already burdened transactional systems. Database tuning for reports is substantially different than for inserts, updates, and deletes (transactions). The type of queries, the volume of the data, and the number of users all add to the processing complexity. Ultimately and invariably, design decisions are made that compromise transaction processing, report processing, or both. Complex reporting must be off-loaded from transactional systems. Techniques include straight replication, the

creation of operational reporting marts, and the integration of transactional data into an operational data store. If the same data has a requirement for historical analysis, then the enterprise data warehouse shall be used.

Analytical Reporting against Historic Data

When historical data (defined as the history of changes to a data record, not the history of transactions attached to a current record) is required for analysis, it shall be provided through the enterprise data warehouse environment. An example of a historical change to reference data would be the change of the name of Washington Township to Robbinsville Township in 2007. It is important to be able to report on all records that occurred in the municipality regardless of name, but it is also important to know what the name was at the time of a particular transaction.

Other types of data exist in the form of snapshots (data that reflects a moment in time, such as a balance sheet), and versions (data that represents the different versions of a record, such as an employee). These data formats are typically not managed in transactional systems. New Jersey manages this data in the enterprise data warehouse in the form of slowly changing dimensions, snapshot fact tables, and profiles. This provides the historical context for reference data.

New Jersey Data Stores

A data store is any database or data repository. Different data stores serve different purposes, and the purpose is independent of the database or repository technology employed. The following specialized data stores are part of the NJSDI and are consistent with New Jersey's IA design patterns.

Transactional Processing Source Systems

These data stores are where the results of business transactions with the State or events of interest to the State are stored. They can be in relational, hierarchical, or file-based database management systems. They can be on a mainframe or on a distributed (network) server. They can be batch processing systems, on-line transactional processing (OLTP) systems, or a hybrid.

Operational Data Store (ODS)

An ODS is a central repository of current operational data initially gathered from a variety of existing transactional systems to present a single rational view of operational data for a single subject area or business unit, or for an entire agency or line-of-business group. History should not be managed or stored in the ODS. Some reporting can occur directly against an ODS, but data can also be replicated into operational reporting areas called Operational Data Marts (Opera Marts).

New Jersey Universal Data Store (NJUDS)

The NJUDS is the central repository of Tier 0 (universal) data and Tier 1 and Tier 2 reference data on behalf of the enterprise. It contains published versions of master reference data (such as the table of counties), standard entities (such as the master address file), and conforming data warehouse dimensions (such as the employee profile). The NJUDS provides mechanisms for managing the universal data, and publishing it or making it available to systems in a variety of forms and formats.

New Jersey Enterprise Data Warehouse (NJEDW)

The NJEDW is a central repository of historical data that is gathered from a variety of sources to support data integration efforts. An EDW publishes the single version of the truth that supplies historical data to data reusability partners, as well as to analysis areas called Data Marts. It is not a single database, but a consistent data integration environment that consists of multiple subject areas, staging, archiving and persistent storage and multiple physical databases. It is rarely accessed directly by end-users.

New Jersey's information architecture does not support the development of independent data marts (directly built from source systems). Instead, data should be persisted in the EDW for future use. Data is stored in the EDW in one of several ways: in the form of a fully normalized data model for the subject area, as a persistent file en route to a reporting area, as a historical dimension table (reference table with history), as a snapshot table (event table with history), or as a detailed or summarized fact table (array of measure created from the transactional data). Our EDW environment accommodates data for individual subject areas, agencies, and the State as a whole.

Data Mart

A data mart consistent with the New Jersey IA is a pre-defined and pre-formatted subset of data sourced from the EDW or an Operational Data Store that has been identified based on the questions that need to be answered by the

report community. Data marts are built for the needs of a specific report community, so the same data may exist in many ways and many combinations in different data marts. They may be logical, consisting of views of enterprise data warehouse data, or physical, consisting of extracts of enterprise data warehouse data. Data is represented in a data mart in one of several ways: in the form provided by the transactional system, as a historical dimension table (reference table with history), as a snapshot table (event table with history), or as a detailed or summarized fact table (an array of measures created from the transactional data).

Dependent data marts always receive data from a consistent, integrated source – never directly from individual operational systems – so the answer to the same question from any data mart is always the same. The NJSDI supports the development of dependent data marts (sourced from the NJEDW environment or an ODS) using conforming dimensions (common reference data used by multiple data marts).

NJ Reusable Data Infrastructure Standard and Supported Technologies

Business Intelligence Publishing Tools

These are query and reporting tools that provide rapid development of reports and can be produced by most business people due to a friendly, graphical interface and a semantic layer than hides the complexity of data relationships from report consumers.

The State supports SAP's BusinessObjects WEBI as its enterprise Business Intelligence Publishing Platform. Other reporting platforms in use include WebFocus and CrystalReports for ubiquitous reporting.

Extract, Transform and Load (ETL) Tools

ETL tools are used to move and transform thousands of records in a bulk fashion and are designed and administered in a graphical environment. These tools learn about data and systems and enable reuse of knowledge on subsequent projects.

The State's ETL Platform is IBM's InfoSphere DataStage, which is web services-capable, XML-aware enterprise integration platform that supports both high volume batch integration and individual transaction integration in real time.

Enterprise Application Integration (EAI) Tools

EAI tools are used to integrate common data across multiple systems at the transaction level, reusing information quality data (metadata). The State requires XML-based web services in a services-oriented architecture (SOA) framework for transaction-level integration.

The State's supported EAI platforms include IBM's InfoSphere DataStage with RealTime Services and WebSphere Message Broker and Enterprise Service Bus.

Metadata Management

The New Jersey IA requires management of metadata, or information resource data, which can include such diverse categories as data dictionaries, data models, process rules, data lineage, system documentation, transformation rules and security information. Metadata management tools share definitions of data between each other and the systems that they document. When possible, common data names and definitions are shared between systems.

The State's standard data warehouse metadata manager is IBM's InfoSphere Metadata Workbench. The State's standard metadata catalog and reference data repository is housed on the Socrata's open government platform. Metadata collection is model-driven using CA's ERWin modeling platform.

Data Modeling

Data modeling tools are used to document, locate and reuse data as well as to describe the relationships between data and systems.

The State uses a number of data modeling tools, such as CA's ERWin, IBM's Rational Architect, Oracle's Designer, and Sybase's PowerDesigner. The OIT Data Architecture unit uses CA's ERWin for logical and physical modeling of transactional and dimensional systems.

Data Profiling

Data profiling tools are used to discover, document and analyze legacy data, capture metadata, map transformations, and describe the relationships between data and systems.

The State's standard data profiling platform is IBM's InfoSphere ProfileStage.

Data Quality and Cleansing Tools

These tools are used to analyze data values, ensure that data elements are captured and stored in a way to best comply with their business rules and intended application, find patterns of poor quality, standardize addresses, add geographic coding information to records, and perform sophisticated matching of free-form data to find exact or like matches.

The State's standard data quality platform is IBM's InfoSphere QualityStage suite.

Data Mining

Data mining is a sophisticated statistical analysis of data for patterns and clusters. It is not the ability to perform ad hoc queries against data, which is provided by business intelligence tools. Data mining tools can learn from earlier analyses and can look for patterns without guidance.

The State does not have a data mining standard.

Supported Database Management Systems (DBMS) Platforms

The State requires that all new transactional database development be built in ANSI SQL-compatible relational database management systems (RDBMS).

The strategic RDBMS products for the State are Oracle's Enterprise Edition and Microsoft's SQL Server. While the State is researching open source products such as MySQL and PostGreSQL, they are not part of the State's strategic direction at this time. The State is also researching the use of NoSQL platforms such as HADOOP.

The State maintains the following mainframe legacy databases: Bull's DMIV, CA's Datacom, IBM's DB2, IBM's IMS, and Software AG's Adabas. The State does not anticipate significant new development taking place on any of these platforms, and is engaged in various initiatives to phase out these environments.

The State maintains a variety of flat file management systems with a strong emphasis on IBM's VSAM for non-DBMS legacy applications, as well as a legacy environment of Focus files. The State is migrating its Focus solution to a data warehousing environment built with Oracle and various Business Intelligence Publishing tools.

Application Development and Infrastructure

NJ IT Architecture Stack

Presentation & Portal Services		
Integration & Messaging		
Application Development & Infrastructure		
Data Management		
Enterprise Servers and Operating Systems		
Next Generation Services Network		
Facilities and Environmentals		

(Side labels: 24x7 Enterprise Help Desk; 24x7 Enterprise Systems Management; Identity Management; Performance Assessment)

The strategic environments for new applications are service-oriented designs using Java J2EE components running primarily on Oracle Java Enterprise System application servers and Microsoft .Net components running on Intel platform servers. All programs should be designed with the goal of developing reusable components. The benefits of building reusable components are evolving into an enterprise framework where common functionality can be shared across applications and platforms. Authentication and authorization should be designed using the New Jersey Identity and Access Management Infrastructure currently provided by the myNewJersey Portal, which leverages pre-defined communities of users and applies role-based policy against those user communities.

J2EE Application Hosting Environment

The State's primary J2EE hosting environment is based on the Oracle (Sun Java Enterprise System) Application Server 9, also known as Glassfish, Enterprise Edition, which has been implemented in standalone as well as clustered configurations. J2EE application design, dependent upon security requirements, usually conforms to a multi-tier architecture as depicted below:
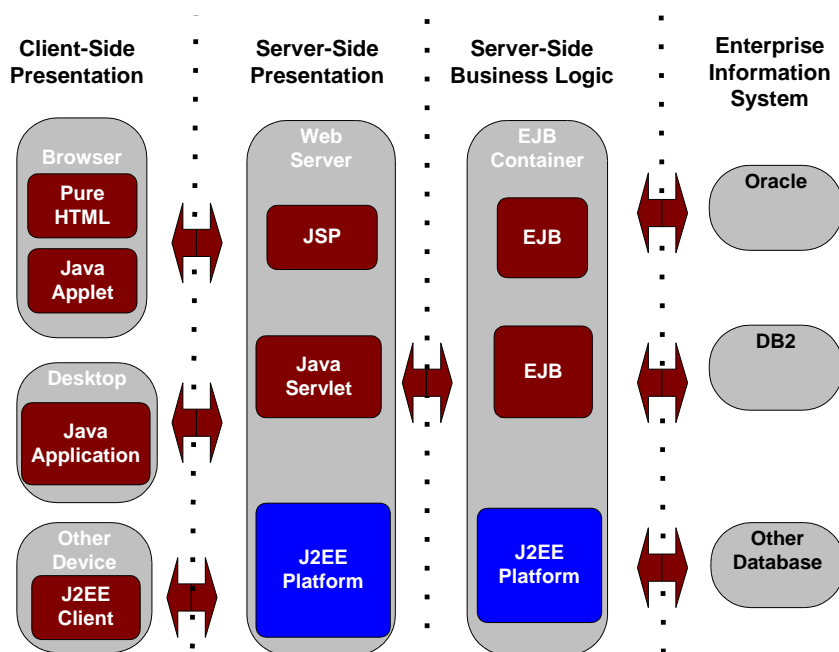
Among the key architectural elements are:

Core Functionality

- Certified compliance with J2EE 1.4 (J2SE 1.4, EJB 2.0, JDBC 2.0, Java Servlet 2.3, JSP 1.2, JMS 1.0, Java Naming and Directory Interface (JNDI) 1.2, Java Transaction API (JTA) 1.0, JavaMail 1.2, Java Activation Framework (JAF) 1.0, JAXP 1.1, J2EE Connector Architecture 1.0, Java Authentication and Authorization Service (JAAS) 1.0)
- An integrated Java Web Services Pack, including JAXM, JAXP, JAXR, and JAX-RPC
- Enabling existing applications to become new Web services through integrated support of SOAP and WSDL
- J2EE Connector Architecture service provider interfaces
- High-performance Java Message Service (JMS) provider
- Java Transaction Service (JTS) with two-phase commit for managing database services from the leading RDBMS vendors
- High-performance HTTP Server with SSL security, delivering high performance through an advanced multiprocessing, multithreaded architecture; efficient use of kernel threads; and sophisticated memory management
- Server-side HTML (SHTML) and chunked encoding which enhance performance of dynamic content
- Various security standards: SSLv2, SSLv3, Transport Layer Security (TLS) 1.0, X.509 certificates, PKCS #11, FIPS-140, 168-bit step-up certificates
- High-performance container-managed persistence (CMP) engine that supports object-to-relational (O/R) mapping

High Availability

- Separate Business Logic and Persistence Tiers. This enables greater scalability across both the business logic and persistence tier while allowing for integrated installation and administration.
- Distributed, Replicated State Information. Application session state data is automatically replicated and distributed across multiple servers. Any individual component can fail without affecting an application's ability to retrieve the session state.
- Inherent Data Availability. The inherent high-availability features delivered with the integrated HADB (high availability database) offer near-continuous availability for application session state data. Application session state data is synchronously replicated.
- Horizontal Scalability. As the load and throughput requirements grow, additional servers for application support and session state maintenance can be easily added without downtime - yielding near linear horizontal scaling.
- Self-Repair. High-availability technology identifies failed servers and can automatically repair to alternative servers, raising overall system availability.
- Shared-Nothing Architecture. The underlying architecture used by Oracle's high-availability technology is inherently distributed, eliminating bottlenecks and facilitating high throughput across multiple servers.
- "Five 9s" availability for Application Server session state persistence.
- Uninterrupted services by providing online upgrades of both software and hardware for better serviceability.

.Net Application Hosting Environment

The Microsoft .Net environment is also built using a multi-tier architecture implementing a web services approach using C#, COBOL and Visual Basic.

Client-Side Presentation	Server-Side Presentation	Server-Side Business Logic	Enterprise Information System
Browser — Pure HTML	**Web Server** — ASP.Net Web Form, ASP.Net Web Form, DOT Net Platform	**App Server** — SOAP Web Service, SOAP Web Service, Class Lib, DOT Net Platform	Oracle, SQL Server, Other Database
Desktop — Windows (Console) Application			
Other Device — Other Application			

Core Functionality

- .Net framework 1.1 and 2.0, which contains Common Language Runtime (CLT) and a collection of .Net application classes
- Internet Information Server 6 (IIS6) is used to host web applications and web services
- Application Center 2000 SP2 (AC2000) is used to control web application deployment, network load balancing and component monitoring
- Standard protocols: SOAP, XML, WSDL and Universal Description, Discovery and Integration (UDDI); ASP.Net
- Database connectivity to Oracle and Microsoft SQL Server, ODBC, OLE DB and XML data sources
- Authentication protocols: Basic, Digest, NTLM, Kerberos and SSL/TLS client
- Cryptography features for encryption, digital signatures, hashing and random number generation including DES, Triple DES, RC2, RSA, DSA, XML digital signature specification, and hashes (MD5, SHA1)

High Availability

- Separate Business Logic and Persistence Tiers. This enables greater scalability across both the business logic and persistence tier while allowing for integrated installation and administration
- Drive Redundancy. Each server contains two mirrored drives and a hot spare which allows the server to continue functioning even if two drives are lost
- Server Redundancy. There are duplicate servers in both the public and secure tiers to enable workload balancing and continuous availability in the event of a server failure
- Horizontal Scalability. As the workload increases, additional servers for application and web support can be easily added
- Clustered Servers. The web and app servers are clustered using Application Center 2000 which provides load balancing, failover support and monitoring capabilities
- Network Load Balancing. Cisco switches distribute work across web servers
- Deployment Servers automate application change management

- Tivoli tools are used to monitor the health of servers to detect and correct problems

eForms

The State has implemented an eForms platform composed of the Adobe Forms Server version 7 using the LiveCycle Forms 7.0 product.

This eForms platform provides electronic forms to New Jersey's internal and external users quickly and efficiently by delivering an XML-based form as a PDF or HTML formatted page to any browser on any device without the need for a download or plug-in. Users with Adobe Reader 7.0 or higher will also have the ability to work offline and submit the form electronically when it has been completed.

Document Management

The State of New Jersey has in place substantial resources and operations for the processing and management of electronic documents.

Automated document management/storage systems include, but are not limited to, systems based on electronic work flow automation, on-line storage and retrieval of record images, Internet-based filing/record retrieval, electronic payment systems (i.e., electronic fund transfer (EFT), e-check, , etc.), email archive systems and records management systems or combinations using technological platforms such as these. New Jersey Enterprise Services include mail processing, remittance processing, document screening/preparation, electronic scanning, index/application data capture, and hosting of electronic images on server platforms.

In virtually all new systems there are potential elements for document management functions. Agencies should seek to utilize existing State document management services as a first choice rather than acquiring or building duplicative services models.

Instruction to Agencies

Agencies are to conform to Statewide Information Technology Strategic planning processes as outlined and administered by the Office of Information Technology (OIT). By Circular Letter 07-11-OMB, agencies are required to review all planned major enhancements to existing systems and new initiatives with the State's Automated Records Management Systems Committee (ARMS) for opportunities to leverage existing State operations as part of solutions which may include elements of records management.

To take advantage of existing Enterprise Services, agencies should contact the Automated Records Management System (ARMS) Committee as early as possible in an initiative's life cycle. This Committee will assist agencies with their plans for new or enhanced information processing systems where they may be related to or may take advantage of existing Enterprise Services to perform all or part of document management processes. Early notification and dialogue with the ARMS Committee regarding planned systems and services will greatly facilitate the review and approval process.

Instruction to Vendors

Vendors working with the State must review and consider the State's capabilities regarding document management services when proposing solutions to agency needs, requests and Requests for Proposals (RFP). Where possible the State will seek to leverage existing facilities and document management processes and services in conjunction with new initiatives.

Automated Records Management Systems Committee (ARMS)

Circular Letter 07-11-OMB (C.L.) establishes a central, inter-agency committee called the Automated Records Management Systems Committee (ARMS), which consists of representatives from the Divisions of Archives and Records Management, and Revenue – Strategic Document Services; and the Office of Information Technology. ARMS is responsible to coordinate the use automated records management and storage systems and policies within the State. These systems and services encompass a broad range of activities – from electronic scanning, indexing and storage of public documents to electronic government applications that supplement or replace paper-based systems.

ARMS seeks to accomplish several strategic goals:

- Comprehensively address State-wide records management and image processing systems/services planning and development, with emphasis on maximizing use of existing in-house facilities;

- Reduce redundant and inefficient system purchases;

- Increase cross-agency sharing of records and information resources;

- Ensure effective use of automated records systems and services on a sustained basis;

- Contribute to the continuing improvement of State government services; and

- Foster adherence to core records management standards; and coordinate information technology and records management planning.

The ARMS Committee can be reached via the Chief Technology Officer (CTO) of OIT, at PMO@OIT.state.nj.us or mailed to OIT -- Project Management Office (PMO), PO Box 212 Riverview Plaza Building 300, 1st Floor.

New Jersey Enterprise Service Packages

Several key services are available to the Executive Branch. Detailed descriptions of these services can be found at: http://www.state.nj.us/treasury/revenue/ARMS/armshome.htm

Records Retention Schedules and Requirements

Proposed systems should provide for and adhere to the State's retention schedule requirements. The State's General Records Retention Schedule can be found at: http://www.state.nj.us/state/darm/links/pdf/g100000.pdf

Technology Infrastructure

While the State may have various implementation of vendor software which accomplish scanning and imaging operations, the primary software that is in use is the FileNet product line from IBM. Application integration for scanning and imaging solutions will utilize interfaces into the FileNet software where they are to utilize existing services. For specific details on the infrastructure contact the ARMS Committee.

Legacy and Mainframe Services

The State has Bull and IBM enterprise systems which host applications for the law enforcement, transportation, labor and human services community among many others. Over one million batch jobs and over one billion online transactions are run on these processors each year. The mainframes are geared toward high volume activity and have excellent response time and availability track records. The applications on the enterprise servers can be web enabled.

Geographic Information System (GIS) Services

The State has a goal of spatially enabling any application that would benefit from geo-awareness. The State definition of spatially enabled means that the system is:

- capable of integrating spatial data (e.g., data with a location component) with other business data across multiple, heterogeneous data sources; and
- capable of supporting abstract data types (e.g., images, text, and spatial data), spatial operators and functions, and spatial locator indexes.

Managing and accessing spatial data across the State's IT enterprise is facilitated through a gateway which utilizes a combination of technologies including Environmental Systems Research Institute (ESRI) Arc Spatial Data Engine (ArcSDE). Spatial data is served up in a format that can be accessed by a variety of desktop GIS clients, served out to the Internet using ESRI's ArcIMS and ArcGIS Server technology or by other applications

using standard SQL queries. Spatial data is hosted on an Oracle and IBM AIX platform providing reliability and scalability.

Internet map server technology provides the foundation for distributing high-end geographic information systems (GIS) and mapping services via the Internet. This technology also enables users to integrate local data sources with Internet data sources for display, query, and analysis in a Web browser. We utilize ESRI's Arc Internet Map Server (ArcIMS) and ArcGIS Server. Both are powerful, scalable, standards-based tools used to design and manage web services for map display and geoprocessing. This technology is currently integrated in the State's Shared Server Infrastructure (SSI) using a three-tier application architecture. Legacy applications continue to be supported on ArcIMS; new applications are encouraged to use ArcGIS Server. Both are maintained at a release level at or near the latest available. The open-source GeoServer product is also supported on a limited basis to deliver simple web services for cases where ArcGIS Server is not feasible.

For ArcGIS Server application development, all APIs provided by ArcGIS Server are available in the State's environment, but among the several REST APIs, Javascript is preferred. Silverlight or Flex are available if necessary, but are discouraged because of the limited client platforms supported.

An array of web services is maintained on the shared infrastructure to meet common functional requirements such as address geocoding, reprojection of data between different coordinate systems, and other similar tasks.

Any proposed solution that includes a GIS component and/or incorporates spatial data is evaluated, planned, designed, and implemented in concert with the OIT Office of GIS. Applications that are geo-enabled are in compliance with the OpenGIS Consortium specifications for spatial data (http://www.opengis.org/). The State of New Jersey's preferred GIS software platform is the ESRI set of products and tools (http://www.esri.com/).

File Transfers

The State has multiple methods of secure file transfer.

The State supports multiple secure file transfer methods for data transfers internally within the Garden State Network or external to an extranet or business partner. The transmissions of all files are transferred securely using the State file transfer methodology.

The State of New Jersey provides both asynchronous and synchronous file transfer methodologies.

Synchronous:

- Connect:Direct Secure ++.
- FTPS over SSL.
- SFTP (FTP over SSHv2 or greater).

Asynchronous:

- DataMotion is a supported option for non-automated or "ad-hoc" file exchange. DataMotion-DataBridge is a supported option for automated file exchange.

The State also supports Extranets to transfer data over dedicated lines, Virtual Private Networks and Extranets between the Garden State Network (GSN) and Business Partners. The agency or Business Partner is responsible for all costs associated with this method.

ePayment

OIT maintains a set of enterprise ePayment web services that provide Internet based payment processing to State agency applications. The ePayment web services allow custom developed Web based applications to either process:

- Credit card transactions by interfacing with a payment gateway provider; or
- eCheck transactions by allowing governmental entities to accept electronic checks via the Internet

Implementation of the ePayment module is facilitated through web services protocols. As such, they can be used with any compliant application in the .NET and J2EE environments. Developers can discuss application requirements with the ePayment Administration Staff at epayment_admin@oit.state.nj.us.

Single Sign-On

See section on Identity Management, Authentication & Authorization Services.

Enterprise eMail Services

The Office of Information Technology maintains a highly available, redundant enterprise infrastructure to facilitate inbound and outbound email processing for State agencies. Gateway services include message routing, reputation, anti-virus and anti-spam scanning.

All inbound and outbound emails are scanned at the gateways for malicious content and email source reputation.

The State is in the process of consolidating to one messaging platform – Microsoft Exchange. This consolidation will create a centralized Active Directory Resource Forest to support a statewide messaging and calendar platform based on Exchange Server 2013 including the necessary systems to monitor and manage the new environment.

Microsoft Office 365 is currently being implemented to provide cloud-based email services in a hybrid mode alongside the State's premise-based Microsoft Exchange infrastructure.

Cloud Services / Software as a Service (SaaS)

The State is actively integrating cloud-based services into its architecture:

- SumTotal Systems Maestro (formerly GeoLearning) is the State's enterprise eLearning service for State and Local employee training.

- Salesforce.com is the State's enterprise Customer Relationship Management service for the Government-to-Business vertical.

- Proofpoint provides email reputation filtering.

- VeriSign provides 2^{nd} factor (strong) authentication services via its VIP service.

- DataMotion provides cloud-based file transfer services.

- Microsoft Office 365 is currently being implemented to provide cloud-based email services in a hybrid mode alongside the State's premise-based Microsoft Exchange infrastructure.
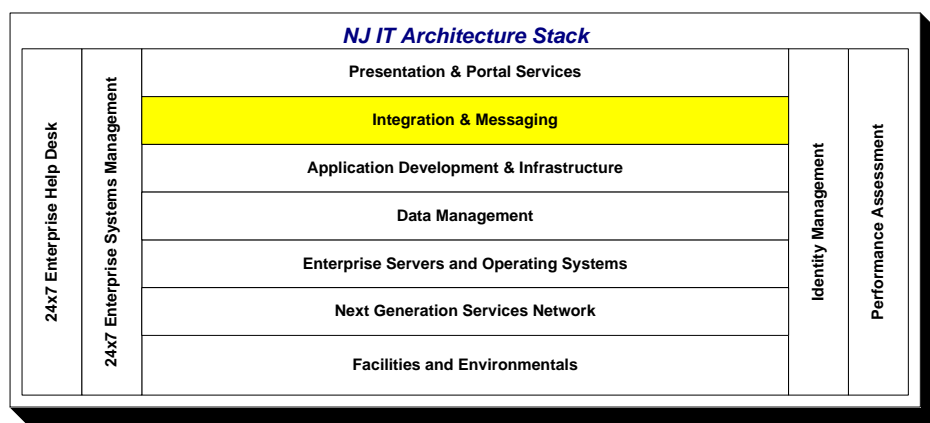
In each implementation, the State has carefully approached SaaS in an integrated fashion, incorporating both data integration at the back end as well as Identity & Access Management integration at the front end.

Integration & Messaging

NJ IT Architecture Stack

Presentation & Portal Services		
Integration & Messaging		
Application Development & Infrastructure		
Data Management		
Enterprise Servers and Operating Systems		
Next Generation Services Network		
Facilities and Environmentals		

(Left vertical columns: 24x7 Enterprise Help Desk; 24x7 Enterprise Systems Management. Right vertical columns: Identity Management; Performance Assessment)

Message Oriented Middleware

The State has implemented IBM Websphere MQ (formerly MQ Series) in many mission critical application environments for enterprise messaging between systems. Websphere MQ is currently in production on the Sun Java Enterprise System Application Server platforms for connectivity to the J2EE application environment.

Enterprise Application Integration (EAI)

An EAI solution enables real-time data and workflow integration from one system to another. The State's Enterprise Data Integration platform, DataStage, when used with the State's message transport standard, Websphere MQ, provides cost-effective real-time application integration to meet many business requirements.

Enterprise Service Bus (ESB)

The State has implemented IBM's Websphere Message Broker at the Enterprise Service Bus (ESB) layer.

Host Application Transformation Services (HATS)

The State has implemented HATS on its IBM platform. These tools provide for rapid development of HTML web based applications using existing CICS applications and native JDBC database connections for data and business logic.

CICS Transaction Gateway

Connectivity to CICS from J2EE applications can be accomplished via the IBM Transaction Gateway. Each instance of the Gateway requires the installation and configuration of a client on the J2EE Application Server platform. On the CICS side, mainframe security and CICS Transactions must be established for the appropriate application(s).

DB2 Connect

Connectivity to DB2 is accomplished via a DB2 Runtime Client, which is installed and configured on the J2EE Application Server platform.

Entire X

Connectivity to legacy Adabas systems from J2EE and .Net applications is facilitated by Entire X Broker connectors.

Presentation & Portal Services
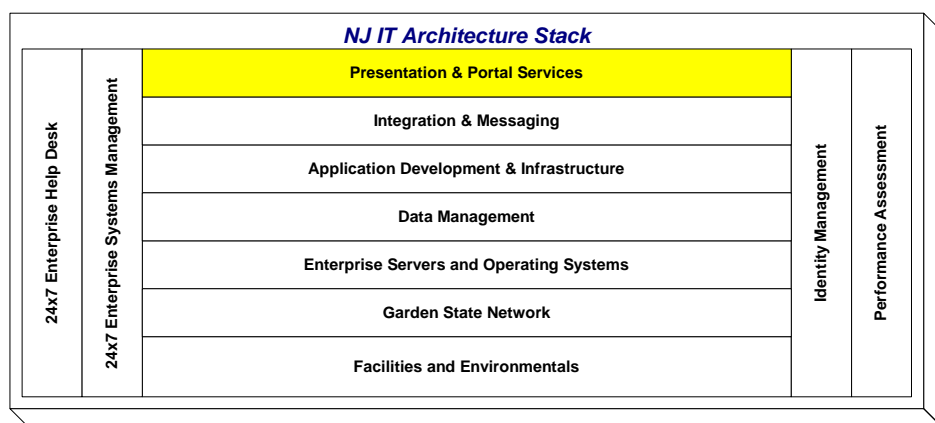
State Portal Overview

The State's Internet Portal provides an identity-enabled array of services including security, user management, single sign-on, personalization, content aggregation and application integration.
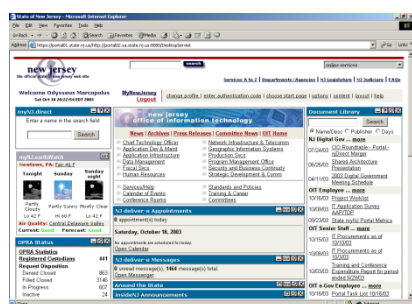
The Portal supports more than 1.2 million registered members across a diverse range of communities – general public, State employees, New Jersey businesses, and local government employees and officials.

One of the larger consumers of the Portal is a Pensions & Benefits self-service for up to one-half million members state-wide including current and retired State and local government employees, teachers, police and firemen.

The Portal infrastructure is based on the Oracle Java Enterprise System Portal Server platform with its internal LDAP directory supplemented by an external Oracle database and custom administration code.

Key features of the Portal infrastructure include:

- Multiple load balanced Web Servers
- SSL encryption of all traffic over the Internet
- On-demand user community creation and management with delegated administration of user policy and access control through an integrated management console
- Dynamic user personalization and customization
- Role based access control (RBAC) with multi-role support, user provisioning, and self-registration
- Delivery of integrated content, applications, and services through customized portlets
- Single sign-on for portal applications
- Integration with existing legacy applications through web services
- Integral Geographic Information Systems engine for location based services
- Rapid deployment of multiple portals for many communities from a single platform architecture

Key collaboration services of the Portal infrastructure include:

- Secure role-based document library that facilitates end-user publishing of materials with email notification to user community
- Secure role-based threaded discussion forum for online collaboration
- Delegated role management with role based email distribution
- Secure role-based content publishing

User services of the Portal infrastructure include:

- End-user self service
- Personalized Weather / Air Quality
- Personalized Events Calendar

Portal User Management

The State Portal provides Role Based Access Control (RBAC) to content and services. It provides single/reduced sign-on capabilities, aggregated content delivery and delegated user management services for online State services. Access control is managed through the assignment of roles via delegated user administration.

Users can "self-register" for access to public web content only. Additional access to secure services can be acquired by either the issuance of an authorization code by a designated role manager, or a self-service application allowing a client to register using personal knowledge based questions. The authorization code process includes formal out-of-band communication between the business process owner and the user.

The State Portal currently uses a combination of LDAP compliant directory services and an Oracle based datastore to manage user authentication, demographic and role assignment data.

The State maintains a set of web services to the Portal user management services allowing custom application developers to leverage these authentication and authorization processes.

Member services and content management are based on the concepts of User, Role and Channel.

User

- Any person, public or private, who is registered with the Portal. A person may self-register with the Portal via the Internet by supplying as little information as a name and email address.

Role

- A role defines a group of users who share sufficient common interests to warrant the creation of a Portal-based user group with access to content and/or transactional systems specifically tailored to those interests.

- Users may be assigned one or more additional roles. Roles provide for a centrally managed user environment and each role has a role manager.

Channel

- A content provider designed to be delivered through the myNewJersey Portal page. Channels are associated with one or more roles.

Web Servers

Anonymous access to the State's static public information is provided through the public access Web servers (www.nj.gov). From there, links are provided to individual agency Web servers.

Currently there are a number of production Web servers. One cluster hosts the State's home page and related flat file information (www.state.nj.us). One cluster supports Microsoft IIS web serving, application serving and data serving through SQL Server. One cluster provides a conduit for the business logic for Java applications bound for the public web server.

The primary web server platform is the Sun Java Enterprise System Enterprise Web Server. It provides the following capabilities to State agency developers:

Web Application Development

- Full compliance for Java Servlet 2.3 and JavaServer Pages (JSP) 1.2 specifications
- Support for NSAPI, CGI, CFML, and PHP
- Built-in Java runtime environment with support for the Java Development Kit (JDK) 1.4x release, object serialization, and the JDBC 3.0 specification, including connection pooling, the Java Naming and Directory Interface 1.1 API, and JavaBeans technology
- Session management service to track information for specific users
- Java technology-based application development across JSP and Java Servlet technologies
- WAR file deployment both from command-line and GUI-based interfaces

- JSP component pre-compilation for faster loading
- Reuse of applications and components that are developed separately
- Standard tag library support, enhancing the user customization of JSP tags
- Fast, in-process, pluggable Java virtual machine (JVM) implementation
- Server-side preprocessing of content using SHTML
- Integration with Java optimization tools
- Web Distributed Authoring and Versioning (WebDAV)
- Netscape Application Program Interface (NSAPI) filter

Reliability and Availability

- High server uptime through multi-processing mode and process monitors
- Unique, shared-session objects to provide failover protection and enable multiprocessing support for Java Servlet extensions on UNIX systems
- Reduced server downtime by rotating logs dynamically
- Intelligent load balancing configuration with Cisco Smart Switch for high availability

Management and Administration

- Dynamic reconfiguration of Web server - without restart
- Integration with Lightweight Directory Access Protocol (LDAP)-based directory servers
- Oracle Java Enterprise System Directory Server management of password policies and user groups down to the site level
- Policy agent integration with the Oracle Java Enterprise System Identity Server
- Command-line interface for HTTP server administration, certificate and key management, and Web application deployment

Performance and Scalability

- High performance through an advanced multiprocessing, multithreaded architecture; efficient use of kernel threads; and sophisticated memory management, Server-side HTML (SHTML) and chunked encoding to enhance the performance of dynamic content
- Multiprocessing mode to increase scalability on multiple CPU machines
- HTTP 1.1 and HTTP compression
- Scalable, keep-alive handling

Security

- Support for SSLv2, SSLv3, TLS 1.0, and X.509 digital certificates
- Support for security-based standards such as PKCS #11, FIPS-140, and 168-bit, step-up certificates
- Centralized, certificate-based security with certificate-to-LDAP mapping
- Administrator setting of SSL parameters for each virtual server
- CGIs to be run as different user IDs
- Single sign-on (SSO) across multiple Web applications (or Java Servlet contexts)

Content Management Services

- Full text and attribute searching of documents through built-in search engine

Web Content Management

Autonomy TeamSite provides enterprise web content management services. TeamSite allows web developers to control the look and feel of the finished pages while allowing non-technical users to provide the content that appears in the final product. Application Infrastructure Services provides the technical support for the infrastructure, and Creative Services provides the development expertise by creating the page templates and is responsible for end user training.

Identity Management

		NJ IT Architecture Stack		
24x7 Enterprise Help Desk	24x7 Enterprise Systems Management	Presentation & Portal Services Integration & Messaging Application Development & Infrastructure Data Management Enterprise Servers and Operating Systems Next Generation Services Network Facilities and Environmentals	Identity Management	Performance Assessment
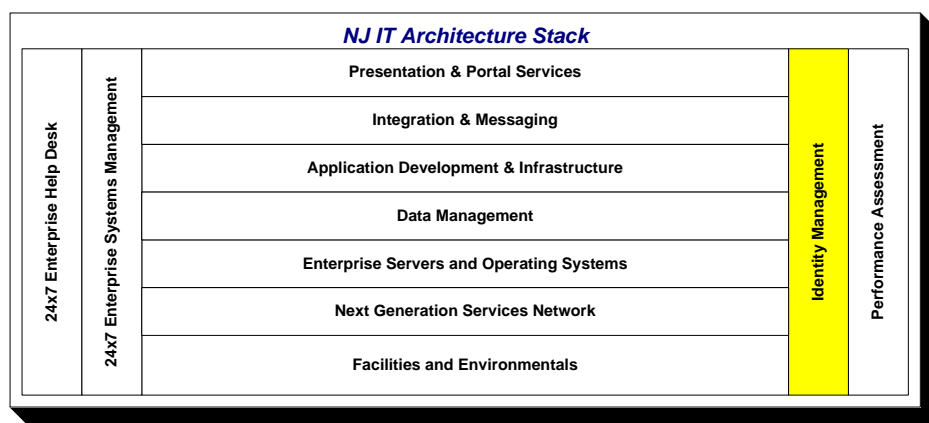
Authentication & Authorization Services

State Internet / Intranet / Extranet Portal

Enterprise Authentication and Authorization services for Internet, Extranet and Intranet applications are currently provided by the State Portal infrastructure. See Presentation & Portal Services for details.

Agent Based Identity Management Infrastructure

The State has implemented an Enterprise Identity and Access Management (IAM) infrastructure to provide a broader array of authentication and access control services. Portal authentication and access control will migrate to this infrastructure, with the Portal becoming a consumer of identity services - as opposed to its current role as provider of identity services.

This infrastructure is based on the Sun Microsystems Java Enterprise System Directory Server, Sun Microsystems Open SSO Access Manager and Sun Microsystems Identity Manager Server and features a comprehensive user provisioning toolset, helping agencies to manage authentication, authorization and access control for the State's business partners, citizens and employees.

IAM provides enhanced delegated user administration for business owners of applications. Multi-factor authentication is supported and available for both Portal applications as well as non-portal applications. The following enterprise-class capabilities are provided:

- Single Sign-On (SSO)
  - Creates a single sign-on session across heterogeneous applications, platforms, and Internet domains
  - Enforces authentication credentials
- Centralized Authorization Services
  - Provides centralized security policy enforcement of user entitlements, leveraging role- and rules based access control
- Federated Identity Support
  - Liberty Alliance Phase 2 (ID-WSF) and SAML 1.1 specifications compliance enable authentication and authorization across federated business networks
  - Provides interoperability across different vendor platforms that provide authentication and authorization services
- Open Architecture and Comprehensive APIs
  - Employs an open, standards-based design to enable high levels of integration and customization
- Enterprise-Class Scalability and Reliability
  - Multiple load-balanced policy servers, policy agents, and directory instances provide high availability and failover capabilities, eliminating any single point of failure

- Real-Time Audit
  - Provides up-to-the-minute auditing of all authentication attempts, authorizations, and changes made.

This infrastructure supports the following industry standards:

- Java Authentication and Authorization Service
- Kerberos
- Liberty Alliance Phase 2 (Identity-based Web Services Framework (ID-WSF))
- Online Certificate Status Protocol (OCSP)
- SAML 1.1 Specification
- SOAP (Simple Object Access Protocol) 1.1
- SPML (Service Provisioning Markup Language)
- SSL (Secure Sockets Layer)
- XML Digital Signature
- XML Encryption
- LDAP version 2 and version 3
- X.509 Digital Certificates

Provisioning

The State has implemented an enterprise Provisioning infrastructure based on the Oracle Java Enterprise System Identity Management platform. Full lifecycle management for the provisioning of digital and non-digital assets has been implemented at the Office of Information Technology and the Motor Vehicle Commission.

Among the deliverables for this implementation are automated user provisioning, account synchronization, auditing & reporting, delegated administration, password management and demonstrable cross platform support.

Application Specific

User authentication and access to applications can also be controlled directly by an application using a custom authentication module and/or access controls embedded in program code or stored at the data layer.

Mainframe

OIT uses a legacy identity and access management product to enable security on the State's mainframe. The IAM is designed to authenticate users and to protect a variety of the plaform resources. The IAM prevents accidental or deliberate modification, corruption, mutilation, deletion, or viral infection of files. With the IAM, access to the system is denied to unauthorized personnel. Any authorized or unauthorized attempt to gain access is logged. System status can be monitored on a continuous basis, and a permanent usage log can be created. The logging feature, besides helping to identify potential intruders, makes it possible to identify and analyze changes and trends in the use of the system. Settings can be changed on a moment's notice, according to current or anticipated changes in the security or business requirements of the organization using the system.

Enterprise Directory Services

The State maintains a Lightweight Directory Access Protocol (LDAP) compliant enterprise directory service for all State employees (NJ Direct). The directory is based on and supports the following industry standards:

- cDSML v2
- LDAP version 2 and 3 RFCs, including RFC 1274, 1558, 1777, 1778, 1959, 2195, 2222, 2247, 2251, 2252, 2253, 2254, 2255, 2256, 2279, 2307, 2377, 2829, 2830, and 3377

Performance Assessment

<table>
<tr><td colspan="5" align="center">***NJ IT Architecture Stack***</td></tr>
<tr>
<td rowspan="7">24x7 Enterprise Help Desk</td>
<td rowspan="7">24x7 Enterprise Systems Management</td>
<td align="center">Presentation & Portal Services</td>
<td rowspan="7">Identity Management</td>
<td rowspan="7">Performance Assessment</td>
</tr>
<tr><td align="center">Integration & Messaging</td></tr>
<tr><td align="center">Application Development & Infrastructure</td></tr>
<tr><td align="center">Data Management</td></tr>
<tr><td align="center">Enterprise Servers and Operating Systems</td></tr>
<tr><td align="center">Next Generation Services Network</td></tr>
<tr><td align="center">Facilities and Environmentals</td></tr>
</table>
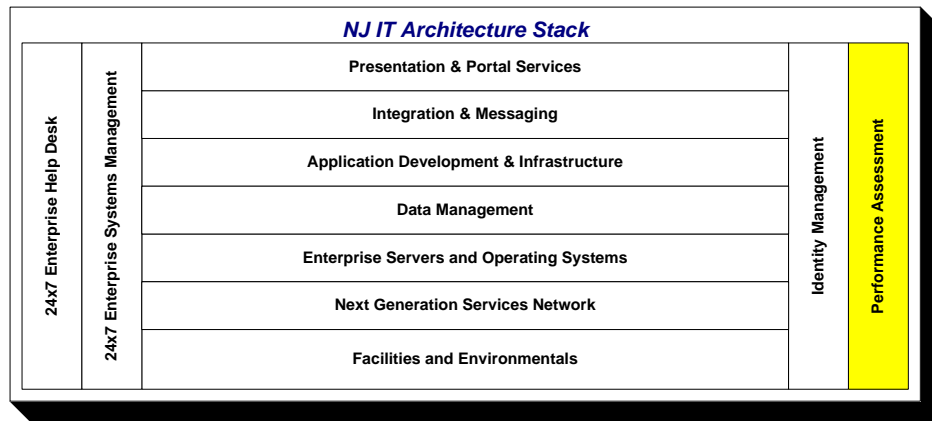
Application Instrumentation and Performance Testing

URL & DB Checker

URL Checker is a State proprietary application that provides monitoring of production browser-based systems. It is a non-invasive performance-monitoring tool that, on a defined schedule, regularly requests expected responses from browser-based systems and provides availability logging as well as technician paging services.

URL Checker is typically implemented on each production J2EE and / or .NET hosting environment in the shared infrastructure. System availability metrics are made available via the State Portal.

DB Checker is a version of URL Checker that is used to monitor production Oracle databases.

HP Performance Center

HP Performance Center is a robust load testing solution that accurately tests the scalability and performance of web applications. The State has implemented Performance Center as an automated software load testing solution to predict how well web applications will handle user load. It is used both during application development and post-deployment to conduct stress testing. Use of this tool has dramatically improved the quality and performance of web based applications. Performance Center is used to create scripts of complex transactions that can then be run in an automated fashion for functional and regression testing of web applications and services.

Bull Mainframe Tools

The Bull environment uses four tools for performance analysis: Video provides information on the jobs that are executing, response times, idle time, and disk and tape usage. Pursue8 displays tape and disk channel usage. Concurrency Monitor displays database conflicts, and Workstation Monitor provides an overview of the workstations that are running and highlights problems.

IBM Mainframe Tools

Omegamon products are used to monitor the operating system, CICS teleprocessing monitor and DB2 database. Trim is used to monitor AG's Adabas database, and Sysview is used to monitor CA's Datacom database.

Network Performance, Application Triage and Performance Service Level Monitoring

The New Jersey Office of Information Technology currently utilizes multiple product sets to monitor, assess, diagnose and provide fault domain or root cause analysis for network and application performance issues. Software and hardware based network and application probes are deployed to perform baseline analysis of existing network environments prior to deploying new applications, upgrading existing applications or implementing new IP services such as VOIP or telephony applications. Application protocols, their respective traffic volumes traversing the local (LAN) and wide area network (WAN) are identified and their bandwidth consumption, average response times, Round Trip Times (RTT) (TCP hand shake), conversation pairs and traffic volumes measured. This analysis can be used as a benchmark comparison against future application or network performance.

Application pre-deployment assessment services are available for departmental and agency clients to assess the network performance characteristics of individual application functions or transactions. Characteristics such as TCP windowing, packet size, conversation flow, node sending and processing behaviors and inspection of processing threads between clients and servers or between nodes in a n-tier hosting environment are examined. The assessment process also includes performance modeling based upon variations in available bandwidth, processing performance and modification in TCP/IP stack parameters.

In the production environment, application triage tools are deployed to monitor and diagnose degradations in application performance and to determine the root cause(s) of poor application performance. Triage tools help to determine whether poor application performance are the result of underpowered hosts, the network infrastructure, the application code or an inefficient host server platform/OS or database.

Pro-active/Real-time monitoring of Application Service Levels is the newest offering available to departmental and agency clients who have deployed application systems within the n-tier environment at the State's data centers. Deployed performance Service Level Management (SLM) toolsets monitor end to end flow of application traffic traversing the E-commerce environment. The SLM process provides both real-time and historical information regarding the performance usage and availability of key business applications (see Appendix 4 –Service Level Management Toolset). The SLM toolsets provide deep dive analysis capabilities of Web based application components including SSL decryption, HTTP page load sequencing, hit level and error analysis, application protocol analysis and decode, including database performance metrics, tiered fault domain isolation analysis, available in real-time or through an extensive set of data mining services.

Network Monitoring

The Cisco Security Monitoring, Analysis, and Response System (CS-MARS) is a Security Threat Mitigation (STM) appliance that monitors the GSN's network health. MARS captures events from reporting devices and evaluates all incidents to determine which default rule will be triggered. The rules that are triggered will determine the resolution of the incident through a threat mitigation process. Through the evaluation process, false positives are determined, consolidated information is distributed through diagrams, charts, queries, and reports.

Vulnerability Management Services

As required by policy and procedures, the Office of Information Technology utilizes vulnerability management as a measure to keep key resources within the Garden State Network safe from hacking and Internet cyber attacks. The Office of Information Technology also oversees vulnerability management efforts in order to ensure New Jersey State Government Executive Branch departments and agencies are meeting policies, regulations, and directives required by New Jersey State Government, the U.S. Federal Government, and private industry. To control and manage risk attributed to security vulnerabilities, the Office of Information Technology provides an Enterprise Vulnerability Management system to departments and agencies. The system is utilized for testing new hardware introduced into network infrastructure and provides an immediate view of network security and compliance posture. The vulnerability management system is also capable of auditing and assessing networks for the possibility of weaknesses that tend to be channels for data and information theft, unauthorized access, or targeted exploitation. Use of the vulnerability management system is guided by the workflow process of detection, removal, testing, and control.

24 x 7 Enterprise Systems Management

NJ IT Architecture Stack

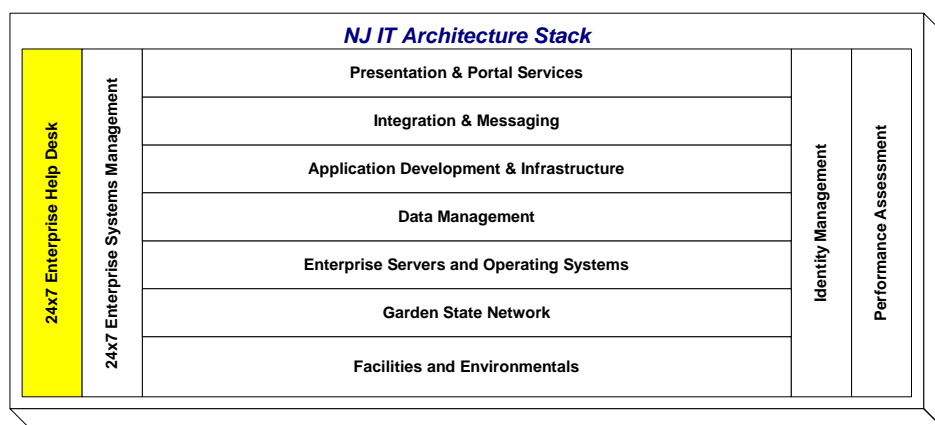| 24x7 Enterprise Help Desk	24x7 Enterprise Systems Management	Presentation & Portal Services	Identity Management	Performance Assessment
		Integration & Messaging		
		Application Development & Infrastructure		
		Data Management		
		Enterprise Servers and Operating Systems		
		Next Generation Services Network		
		Facilities and Environmentals		
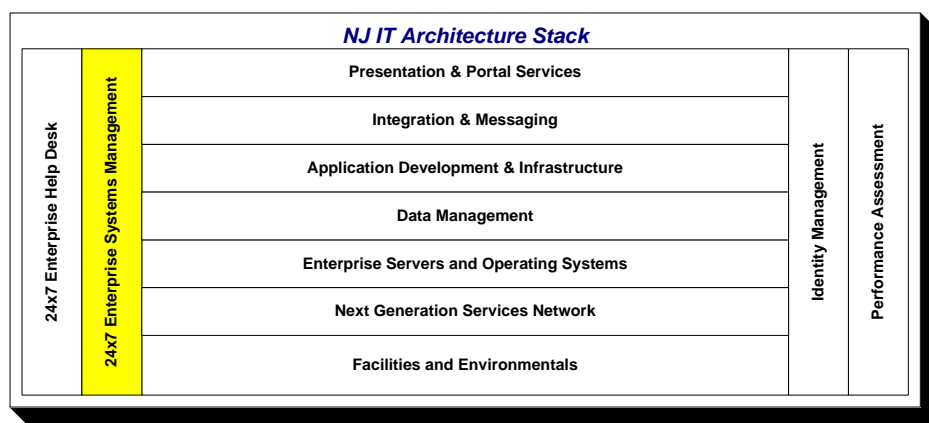
Enterprise Systems Management (ESM) is the proactive monitoring of the New Jersey Shared IT Infrastructure (NJSITI).

ESM at OIT Includes the Following Functionality

- **Network monitoring** provides 7 by 24 monitoring of all Garden State Network devices for Up/Down and hardware problems (see Appendix 5 – Enterprise Systems Management).

- **Server Monitoring** provides 7 by 24 monitoring availability of essential resources and detecting events such as excessive memory or CPU utilization, health status, temperatures, bottlenecks, and application performance through synthetic transactions, etc.

- **Event Management/Correlation** is in many ways the central nervous system of our ESM architecture. Significant events from various monitoring tools are forwarded to the event management software for processing. Through various rule sets and a problem or incident management interface, events considered critical automatically generate problem or incident tickets that are assigned to the responsible groups for resolution.

- **Automated Incident Ticketing/Notification/Escalation** improves client application availability through the automatic notification and escalation of incidents via email and/or sms messages and the integration of problem and change management.

- **Automated Server Application Inventory (ASAI),** an in-house, web-based application, providing a robust hardware and software inventory system of all servers and applications hosted by OIT.

- **Configuration Management** is the detailed recording and updating of information that describes an enterprise's computer system and networks, including all hardware and software components called Configuration Items (CIs).

- **Configuration Management Database (CMDB)** contains all relevant details of each CI and its important relationships, eg. SAN, Network Applications, Locations, Groups, Contacts, etc.

- **Configuration Automation** automatically discovers network devices, servers, operating systems, applications, databases and middleware running in our infrastructure, inventories their configuration settings, and maps discovered components to dependent IT services for change impact analysis. Data and dependency information is captured and presented as snapshots, providing the ability to establish configuration baselines for continuous change tracking and to detect configuration drifts. The Configuration Management Database is automatically updated via Configuration Automation.

- **Presentation (Business View)** includes a front-end Visio-based topology of monitored applications. When a component experiences a state change, the impact of that event is recorded in real-time via a status color change. Drill down capability facilitates root-cause analysis reducing the time to detect and repair.

- **Business Application Performance Monitoring** provides an in-depth view of application service level metrics from real-time end-to-end response time measurements to historical trend analysis data of critical web-based and enterprise applications.

- **Mainframe Monitoring** currently includes monitoring of the operating system, database, and transactional services.

- **Server Backup and Recovery** protects data from hardware failures and other errors by storing backup and archive copies of the data on centralized offline storage. Our Distributed Storage Management solution scales to protect hundreds of computers running a dozen operating systems ranging from laptops to mainframes. Integration with event management is a future consideration.

24 x 7 Enterprise Help Desk

The Enterprise Help Desk / Network Call Center is staffed 24 hours a day, 365 days a year to resolve system outages. The NCC is the first point of contact for users to report IT related problems, request information, request an IT service change, service request, or report a service incident. All calls made to NCC are recorded in the Service Desk Manager System. The system simultaneously e-mails and pages the resources that have been identified to resolve specific incidents. Resources typically include a primary contact, a back-up contact and a supervisor. Resources begin the incident resolution process and update the problem ticket with status information until it is resolved. System users can access this system via a web browser to monitor the resolution status of their incident.
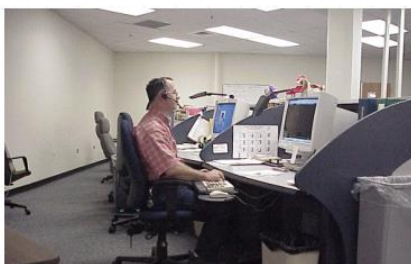
The users can request services with a phone call, an email or by filing an incident request or change ticket via the Service Desk Manager.

The primary objective of the NCC Service Desk is to quickly triage the service requests to either:

- Return the IT service level back to normal as quickly as possible

- Escalate the issue to the next level support

- Provide the desired "how-to" information

- Provide user accounts and password administration

- Provide an update on a previously identified incident or service request

The NCC serves over 20 State agencies on both legacy and new systems. All problems and resolutions are analyzed for performance statistics and problem cause.

Appendix 1 - Products and Technologies

NOTE: This document is not an endorsement of any vendor's products. Vendors who are responding to bid opportunities with the State of New Jersey are not required to propose platforms or products noted in this document unless specifically directed within the requirements section(s) of the bid opportunity.

Category	Product	Support Level*
Application Developer Desktop		
	Windows 7	E
	Windows XP	E
	Windows 2000	E
	Windows 98	S
	Windows 95	S
	Windows NT4	S
Application Development Languages		
	COBOL	E
	C#	L
	HTML	E
	JavaScript	L
	J2EE Java	E
	Natural	S
	Perl	S
	SQL	E
	Visual Basic	S
	XML	E
	.ASP	L
Application Development Tools		
	Adobe	L
	Adobe Creative Suite 6	E
	MS Visual Studio	E
	Oracle Application Express	E
	Oracle Forms/Reports	L
	Quark	L
	Sun Java Studio	E
Application Servers		
	Citrix	L
	IBM Websphere	E
	MS Windows	E
	Oracle	E
	Sun Java Enterprise System	E
	Tomcat	L
Audio / Video		
	Adobe Photoshop CS6	L
	Autodesk Cleaner	L
	Microsoft	L
	Real Media / Windows Media / Adobe Flash / Quick Time	L
Backup and Recovery Tools		
	Cristie Bare Metal Restore	E

Tivoli Suite	E
Symantec NetBackup	E

Business Intelligence (Analysis, Query & Reporting) Tools

Business Objects WebIntelligence	E
Information Builders WebFocus	L
SAS Data Miner	S

Configuration Management

| Automated Server Application Inventory (ASAI) | S |
| CA Configuration Automation and CA CMDB | E |

Customer Relationship Management (SaaS)

| SalesForce.com | E |

Data Integration Tools (ETL, EAI, EII, Messaging, Gateways)

IBM CICS Transaction Gateway	L
IBM DB2 Connect (Gateway)	L
IBM Host Application Transformation Services (Gateway)	L
IBM WebSphere DataStage (ETL)	E
IBM WebSphere Information Services Director (EAI)	L
IBM WebSphere MQ (Messaging)	E
IBM Websphere Message Broker	E
Software AG Entire X (Gateway)	L

Data Management Tools

AppFluent Query Analyzer	L
BI Ready Data Warehouse Mapping Tool	L
CA ERWin (Data Modeling)	E
Data Foundations OneData (Master Data Management)	E
IBM Information Analyzer (Data Profiling Tool)	E
IBM Metadata Workbench (Metadata Repository)	E
IBM Rational Architect (Data Modeling)	L
IBM WebSphere QualityStage (Data Quality Platform)	E
Oracle Designer (Data Modeling)	S
Sybase PowerDesigner (Data Modeling)	S

Data Transfer

Connect:Direct	L
Secure File Transfer	E
Tumbleweed	E
Cyberfusion	E

Database Platforms

Bull IDS2 (Bull DMIV)	S
CA Datacom/DB	L
IBM DB2	L
IBM IMS	S
MS SQLServer	E
Oracle Database	E
Software AG Adabas	L

Directory Services

| Active Directory | E |
| Sun Java Enterprise System LDAP | E |

eForms

Adobe Forms Server	L
Adobe LiveCycle Forms	L
Adobe LiveCycle Workflow	L
Adobe Workflow	L

Enterprise eMail Services

McAfee Anti-Virus	E
Proofpoint Anti-Spam	E

Enterprise Systems Management

CA Service Desk Manager	E
Ground Works/Ngaios	E
ZenOss	E
Axibase Fabrica	E
IBM Tivoli Monitoring (ITM, Netcool/Omnibus)	E

GIS Technology

ESRI: ArcGIS Server – Internet Map/Geoprocessing Server	E
ESRI: ArcIMS– Internet Map Server	S
ESRI: ArcGIS Desktop	E
ESRI: ArcSDE – Spatial Data Hosting	E
ESRI: Metadata Server – Spatial Data Catalog	S
GeoServer – Web Feature/Map Server	S

Groupware Calendar

MS Exchange	E

Groupware Mail

MS Exchange	E

Identity Management / Policy Services

Sun Java Enterprise System Access Manager	E
Sun Java Enterprise System Identity Manager	E

Imaging

FileNet	E

Learning Management (SaaS)

SumTotal Systems (formerly GeoLearning)	E

Legacy and Mainframe Services

CICS	E
TP8 (Bull)	S
VSAM	S

Operating Systems

Bull GCOS8	S
IBM AIX	E
IBM Z/OS	E
LINUX	L
Sun Solaris	E
Windows NT	S
Windows 2000	S
Windows 2003	E

Performance Assessment Tools

Bull: Video, Pursue8, Concurrency Monitor, Workstation Mor	S

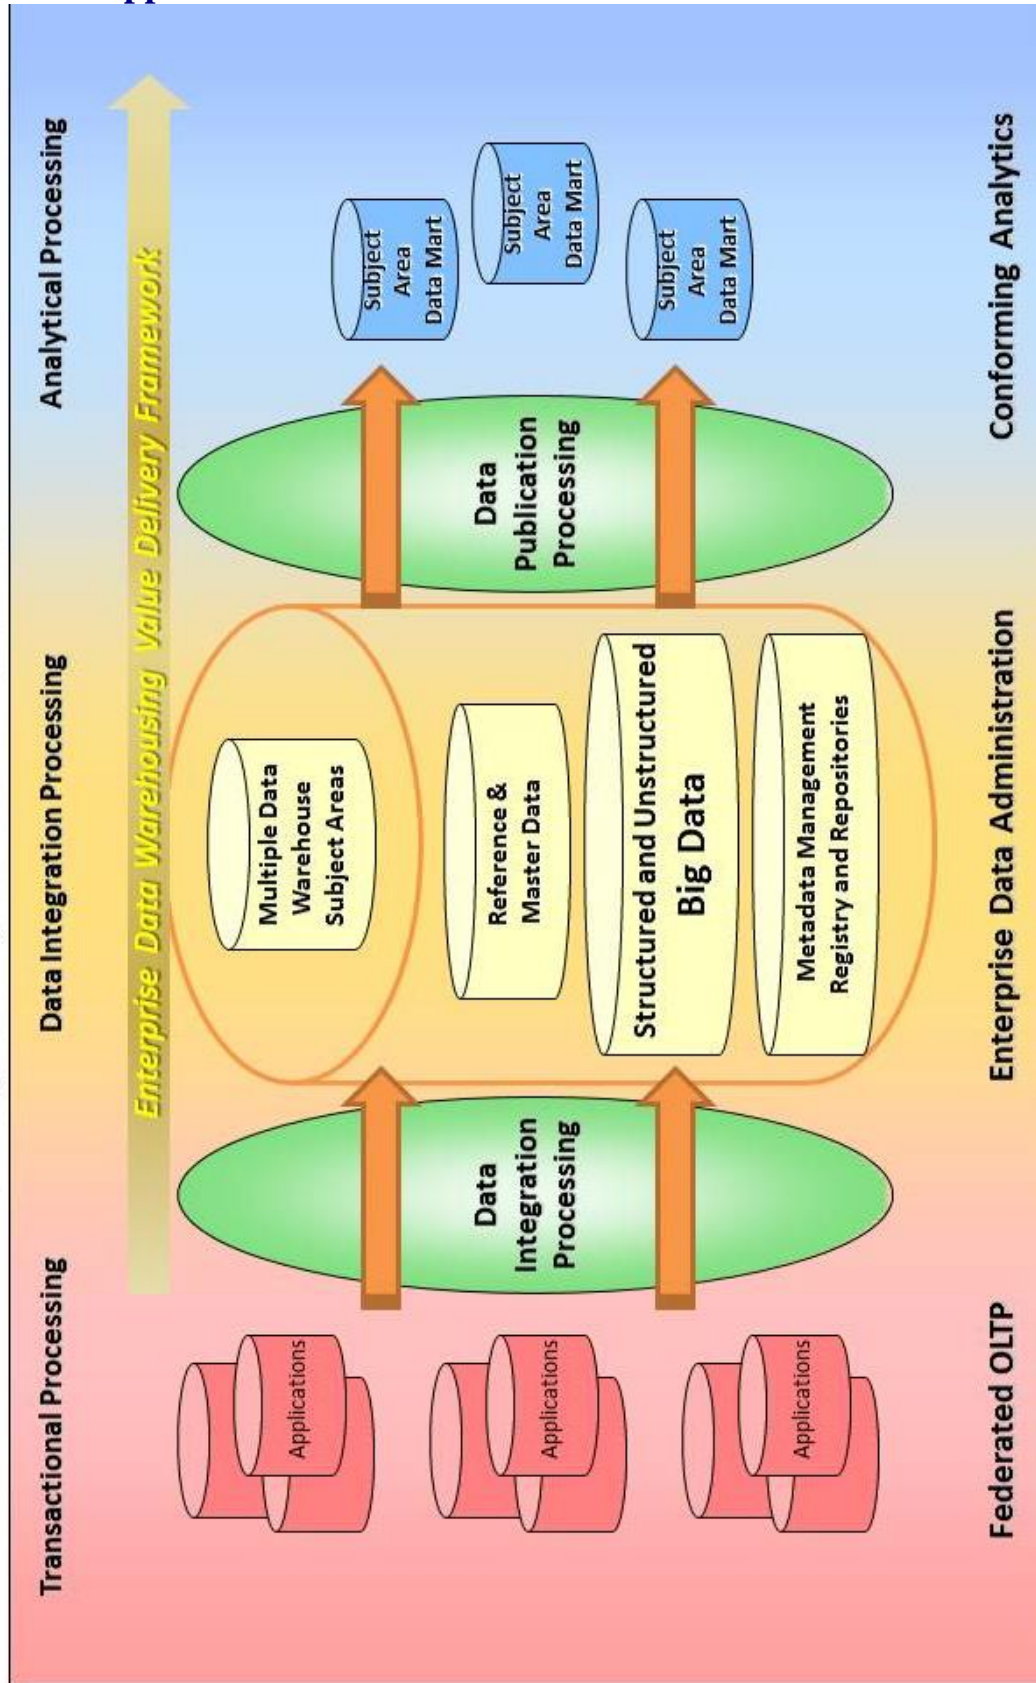|	
IBM: Omegamon, Trim, Sysview	L
LAN/WAN: Compuware Network Vantage, Application Exper	E
Load Testing: HP Performance Center	E

Portal Services

Sun Java Enterprise System Portal Server	E

Print Services

IBM Advanced Function Printing	E
IBM InfoPrint Manager	L
IBM InfoPrint Workflow	L

Security Tools

ACF2	E
SSL	E
VeriSign PKI	E

Software Administration

CA Librarian	E
CVS	E
SourceSafe	L

Transactional System Reporting Tools

Business Objects Crystal Reports	E
Information Builders Focus	S
Information Builders WebFocus	L
Oracle Reports	L

Web Content Management

Autonomy TeamSite	E

Web Servers

IIS	E
Oracle	L
Sun Java Enterprise System	E

* Support Level:

    (E) Enterprise Support
        This represents a technology that is currently supported across multiple State agency initiatives and for which the State has made a substantial investment in infrastructure and staff resources.

    (L) Limited Support
        This represents a technology that is currently supported on behalf of at least one State agency initiative and for which the State has made a limited investment in infrastructure and staff resources.

    (S) Sunset
        This represents a technology the State generally wishes to retire and for which limited or no new investments are being made.
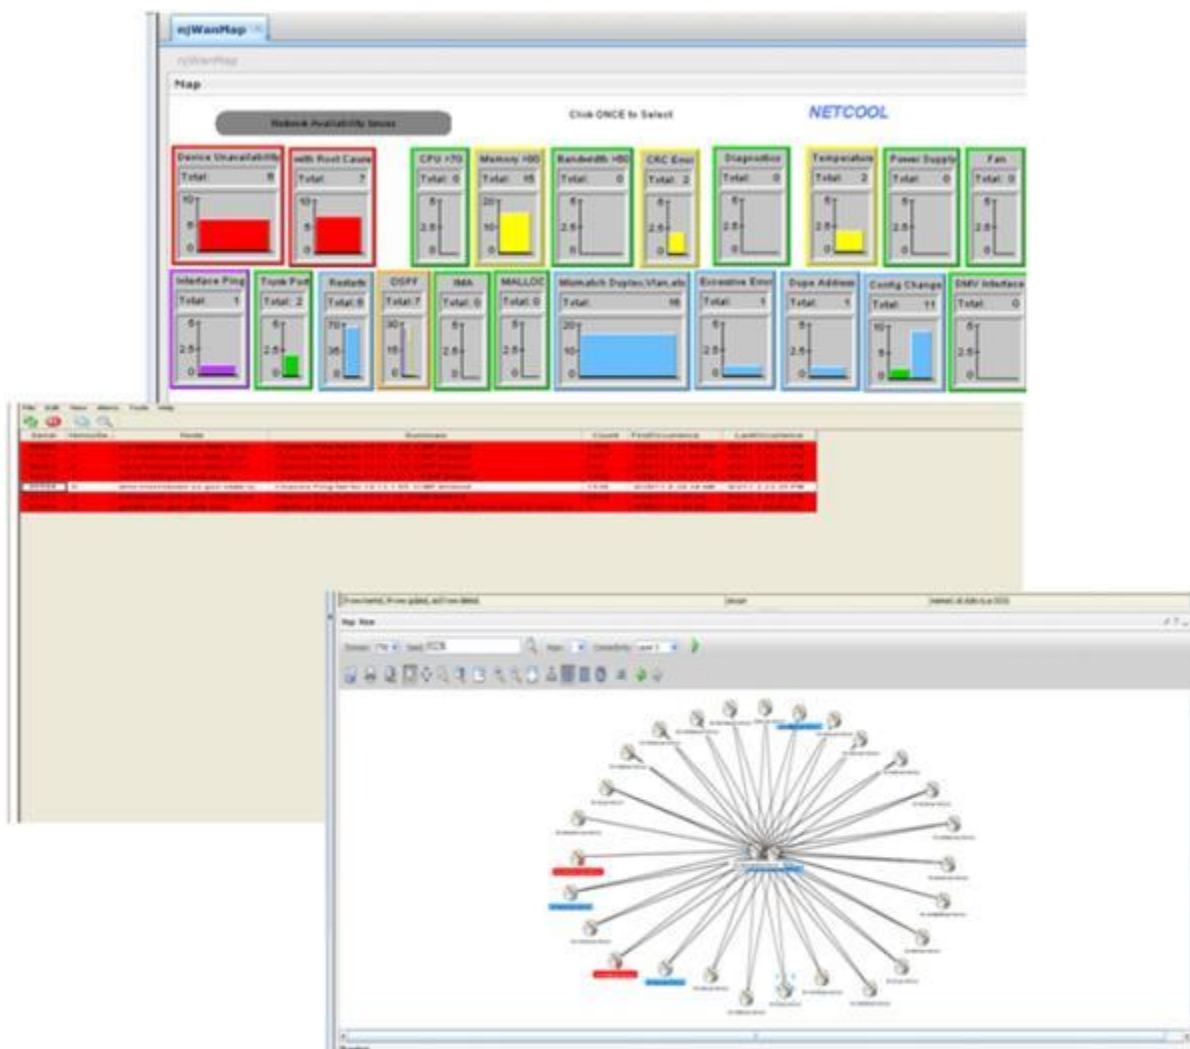
Appendix 2 – NJ Common Information Architecture

New Jersey Data Governance Framework
Conceptual Information Architecture

Transactional Processing — Data Integration Processing — Analytical Processing

Enterprise Data Warehousing Value Delivery Framework

Applications

Data Integration Processing

Multiple Data Warehouse Subject Areas

Reference & Master Data

Structured and Unstructured Big Data

Metadata Management Registry and Repositories

Data Publication Processing

Subject Area Data Mart

Federated OLTP — Enterprise Data Administration — Conforming Analytics

Version 5.01

Revised September 2013

Appendix 3 – Network Systems Management

Appendix 4 – Enterprise Systems Management